

(12) **United States Patent**
Blaner et al.

(10) **Patent No.:** **US 9,208,092 B2**
(45) **Date of Patent:** ***Dec. 8, 2015**

(54) **COHERENT ATTACHED PROCESSOR
PROXY HAVING HYBRID DIRECTORY**

(71) Applicant: **INTERNATIONAL BUSINESS
MACHINES CORPORATION,**
Armonk, NY (US)

(72) Inventors: **Bartholomew Blaner**, Underhill Center,
VT (US); **Michael S. Siegel**, Raleigh,
NC (US); **Jeffrey A. Stuecheli**, Austin,
TX (US); **Charles F. Marino**, Round
Rock, TX (US)

(73) Assignee: **GLOBALFOUNDRIES INC.**, Grand
Cayman (KY)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 116 days.

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **14/035,349**

(22) Filed: **Sep. 24, 2013**

(65) **Prior Publication Data**

US 2014/0379989 A1 Dec. 25, 2014

Related U.S. Application Data

(63) Continuation of application No. 13/921,844, filed on
Jun. 19, 2013.

(51) **Int. Cl.**
G06F 12/08 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 12/0831** (2013.01); **G06F 12/082**
(2013.01); **G06F 2212/621** (2013.01)

(58) **Field of Classification Search**

CPC G06F 12/0831; G06F 12/0822

USPC 711/145, 146, 133, E12.026, E12.069

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,085,295	A *	7/2000	Ekanadham et al.	711/145
6,671,772	B1 *	12/2003	Cousins	711/112
2001/0037426	A1 *	11/2001	Pawlowski et al.	710/260
2008/0086602	A1 *	4/2008	Guthrie et al.	711/144

* cited by examiner

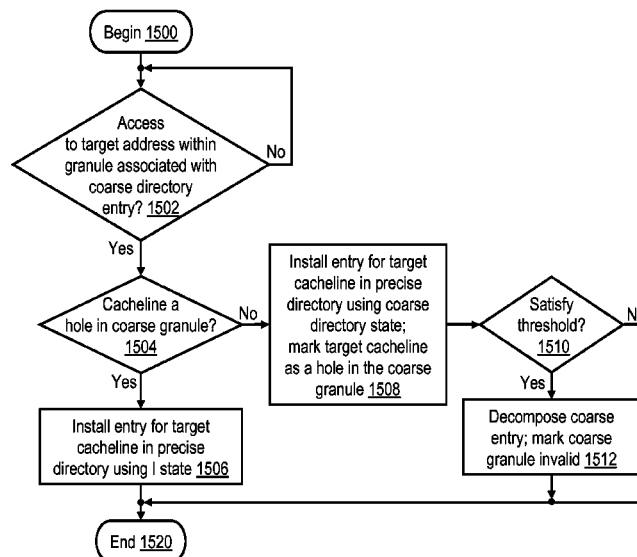
Primary Examiner — Pierre-Michel Bataille

(74) *Attorney, Agent, or Firm* — Russell Ng PLLC

(57) **ABSTRACT**

A coherent attached processor proxy (CAPP) includes trans-
port logic having a first interface configured to support com-
munication with a system fabric of a primary coherent system
and a second interface configured to support communication
with an attached processor (AP) that is external to the primary
coherent system and that includes a cache memory that holds
copies of memory blocks belonging to a coherent address
space of the primary coherent system. The CAPP further
includes one or more master machines that initiate memory
access requests on the system fabric of the primary coherent
system on behalf of the AP, one or more snoop machines that
service requests snooped on the system fabric, and a CAPP
directory having a precise directory having a plurality of
entries each associated with a smaller data granule and a
coarse directory having a plurality of entries each associated
with a larger data granule.

5 Claims, 16 Drawing Sheets






Figure 1




Figure 2




Figure 3




Figure 4




Figure 5

Figure 6

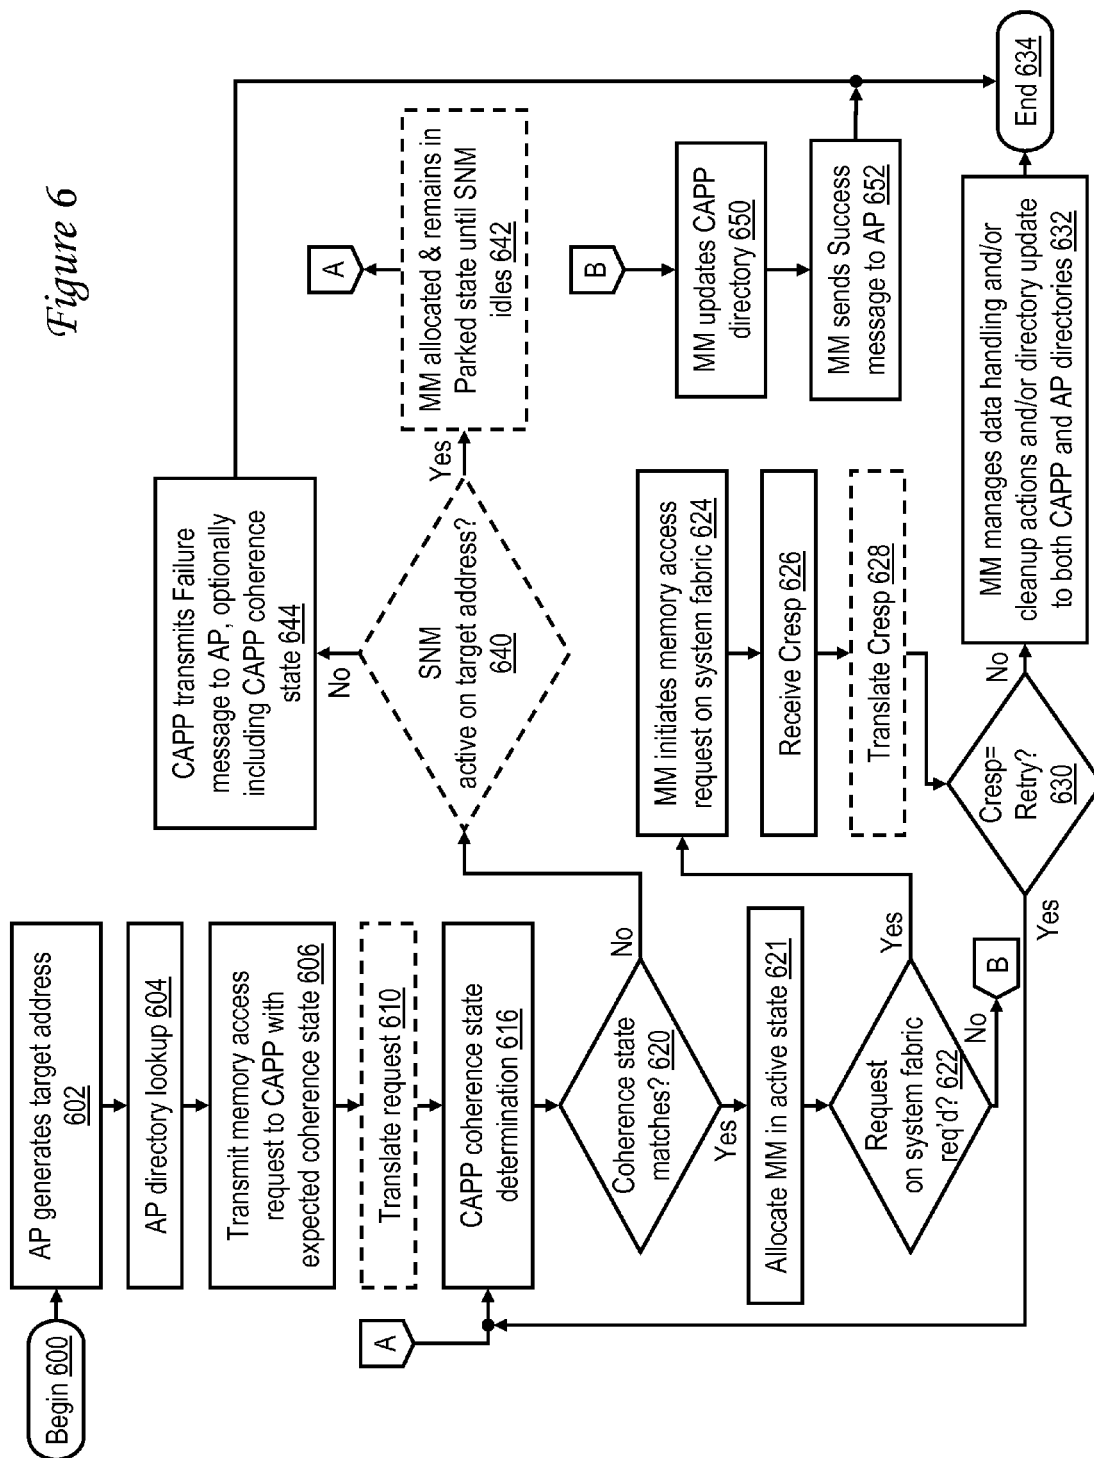
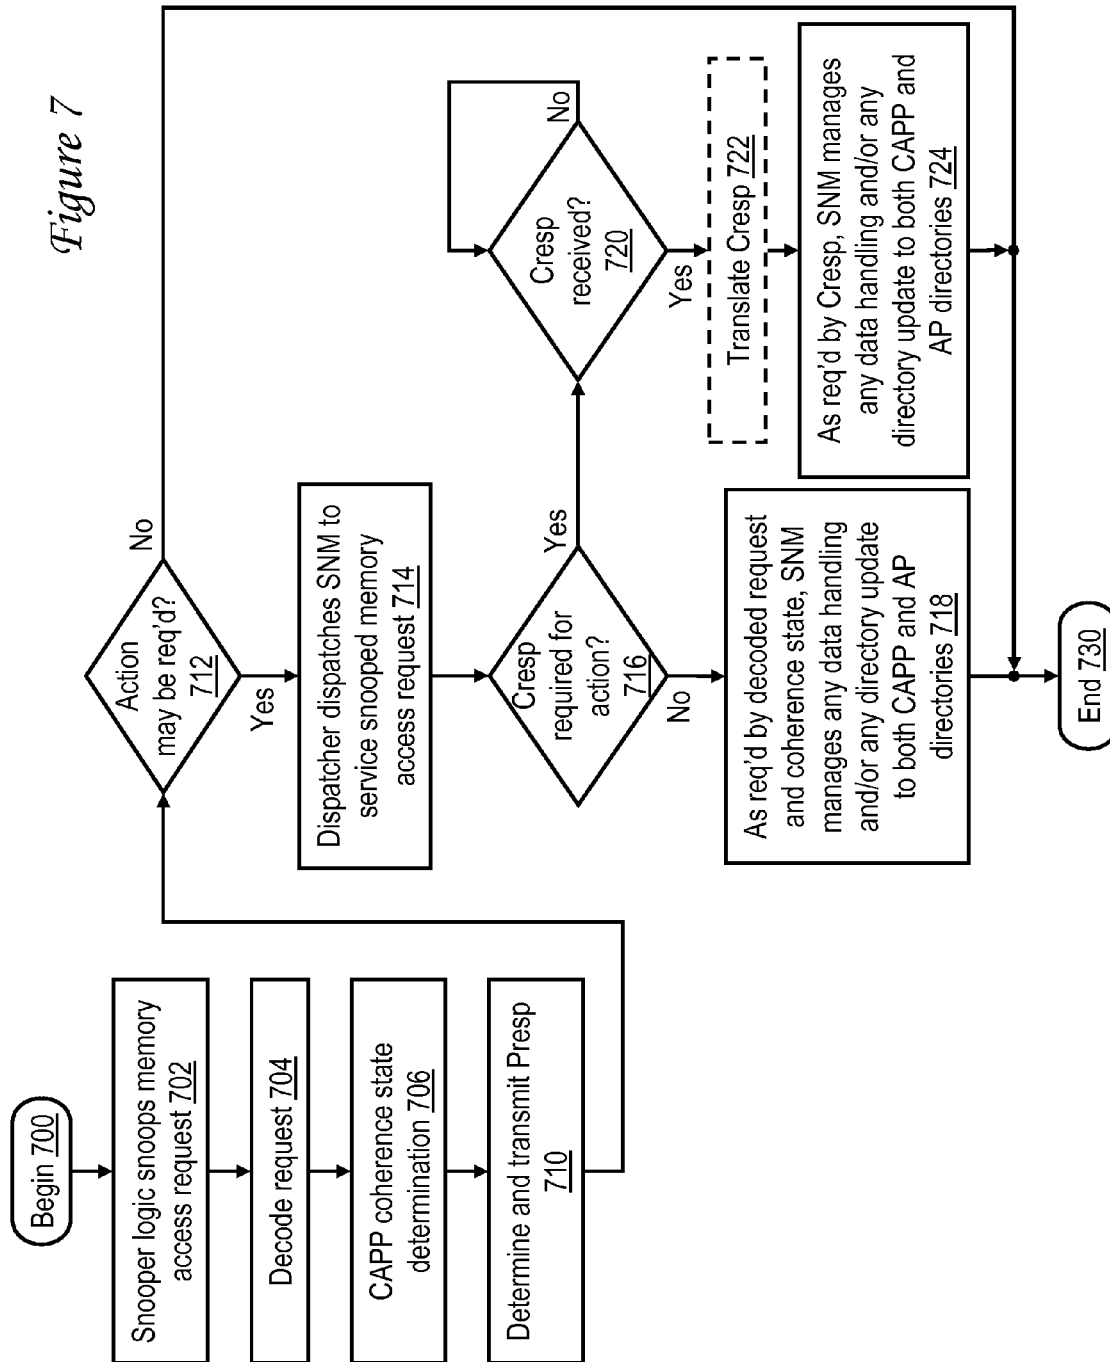


Figure 7






Figure 8




Figure 9




Figure 10


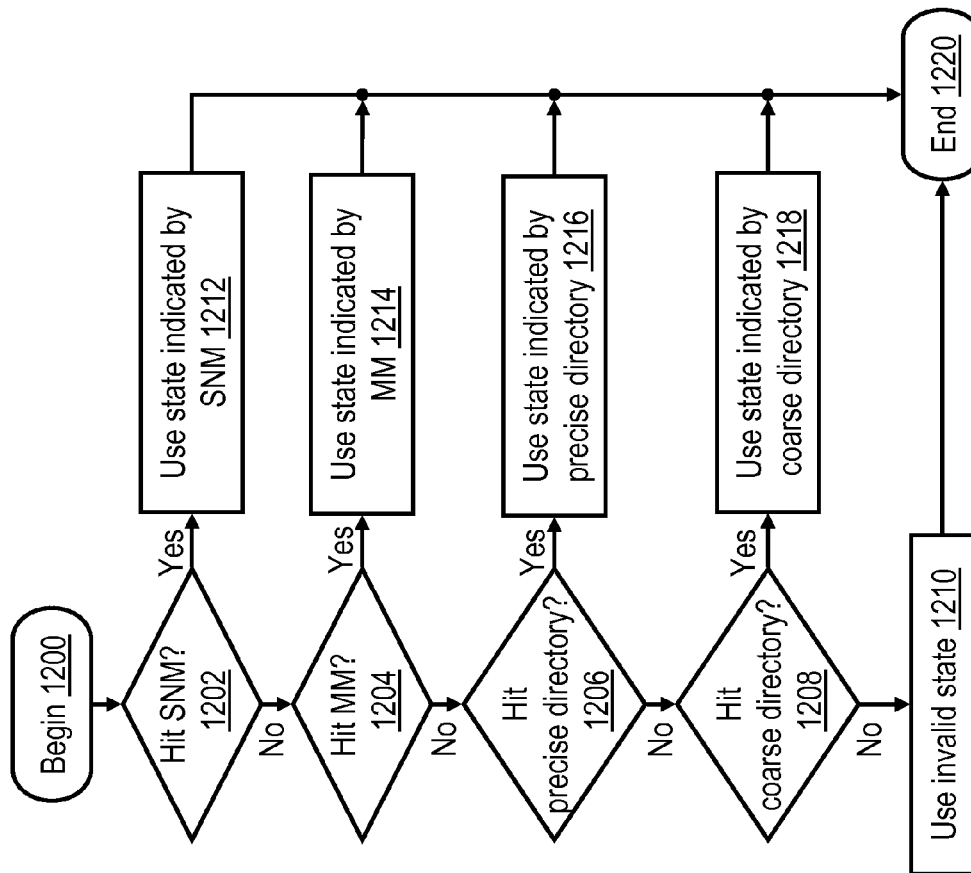
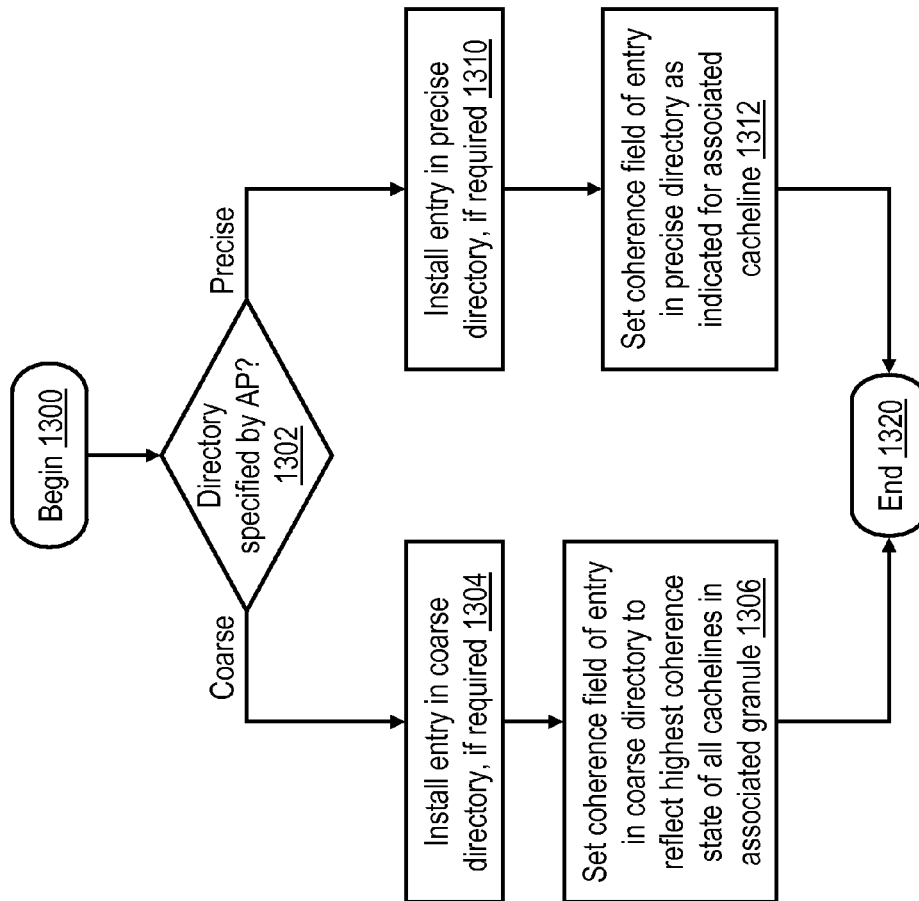
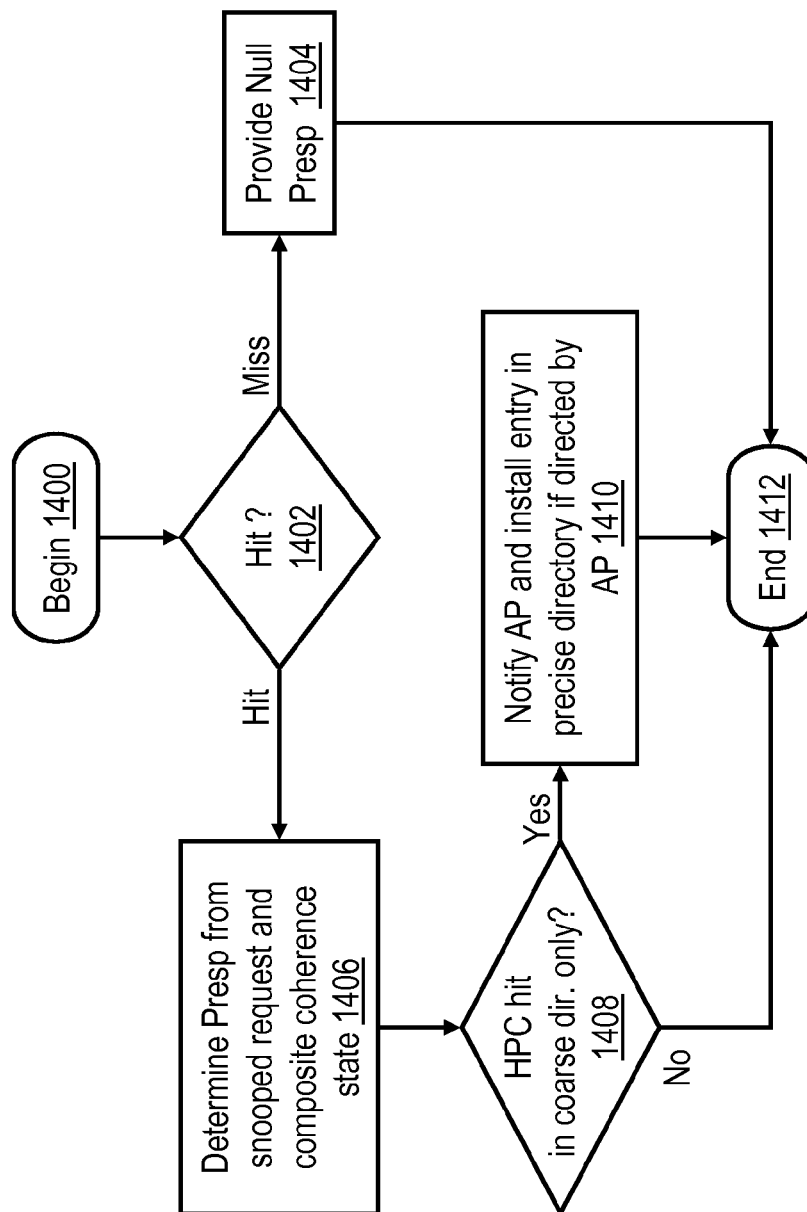
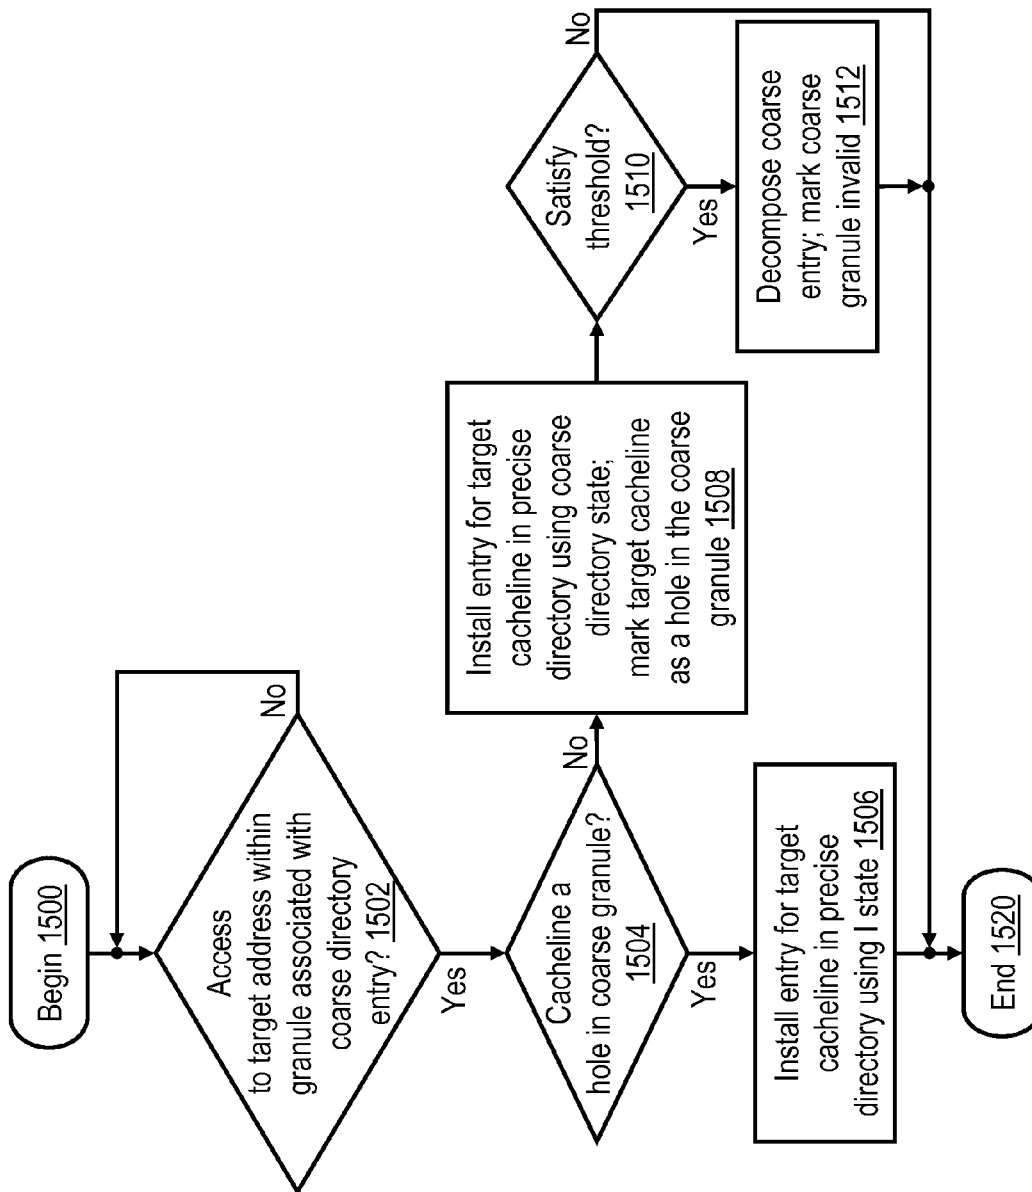


Figure 11

*Figure 12*

*Figure 13*

*Figure 14*

*Figure 15*




Figure 16

1

COHERENT ATTACHED PROCESSOR PROXY HAVING HYBRID DIRECTORY

PRIORITY CLAIM

This application is a continuation of U.S. patent application Ser. No. 13/921,844, entitled "COHERENT ATTACHED PROCESSOR PROXY HAVING HYBRID DIRECTORY," filed on Jun. 19, 2013, the disclosure of which is incorporated herein by reference in its entirety for all purposes.

BACKGROUND OF THE INVENTION

The present invention relates to data processing, and more specifically, to a coherent proxy for an attached processor.

A conventional distributed shared memory computer system, such as a server computer system, includes multiple processing units all coupled to a system interconnect, which typically comprises one or more address, data and control buses. Coupled to the system interconnect is a system memory, which represents the lowest level of volatile memory in the multiprocessor computer system and generally is accessible for read and write access by all processing units. In order to reduce access latency to instructions and data residing in the system memory, each processing unit is typically further supported by a respective multi-level cache hierarchy, the lower level(s) of which may be shared by one or more processor cores.

Because multiple processor cores may request write access to a same memory block (e.g., cache line or sector) and because cached memory blocks that are modified are not immediately synchronized with system memory, the cache hierarchies of multiprocessor computer systems typically implement a cache coherency protocol to ensure at least a minimum required level of coherence among the various processor core's "views" of the contents of system memory. The minimum required level of coherence is determined by the selected memory consistency model, which defines rules for the apparent ordering and visibility of updates to the distributed shared memory. In all memory consistency models in the continuum between weak consistency models and strong consistency models, cache coherency requires, at a minimum, that after a processing unit accesses a copy of a memory block and subsequently accesses an updated copy of the memory block, the processing unit cannot again access the old ("stale") copy of the memory block.

A cache coherency protocol typically defines a set of cache states stored in association with cached copies of memory blocks, as well as the events triggering transitions between the cache states and the cache states to which transitions are made. Coherency protocols can generally be classified as directory-based or snoop-based protocols. In directory-based protocols, a common central directory maintains coherence by controlling accesses to memory blocks by the caches and by updating or invalidating copies of the memory blocks held in the various caches. Snoop-based protocols, on the other hand, implement a distributed design paradigm in which each cache maintains a private directory of its contents, monitors ("snoops") the system interconnect for memory access requests targeting memory blocks held in the cache, and responds to the memory access requests by updating its private directory, and if required, by transmitting coherency message(s) and/or its copy of the memory block.

The cache states of the coherency protocol can include, for example, those of the well-known MESI (Modified, Exclusive, Shared, Invalid) protocol or a variant thereof. The MESI

2

protocol allows a cache line of data to be tagged with one of four states: "M" (Modified), "E" (Exclusive), "S" (Shared), or "I" (Invalid). The Modified state indicates that a memory block is valid only in the cache holding the Modified memory block and that the memory block is not consistent with system memory. The Exclusive state indicates that the associated memory block is consistent with system memory and that the associated cache is the only cache in the data processing system that holds the associated memory block. The Shared state indicates that the associated memory block is resident in the associated cache and possibly one or more other caches and that all of the copies of the memory block are consistent with system memory. Finally, the Invalid state indicates that the data and address tag associated with a coherency granule are both invalid.

BRIEF SUMMARY

In at least one embodiment, a coherent attached processor proxy (CAPP) includes transport logic having a first interface configured to support communication with a system fabric of a primary coherent system and a second interface configured to support communication with an attached processor (AP) that is external to the primary coherent system and that includes a cache memory that holds copies of memory blocks belonging to a coherent address space of the primary coherent system. The CAPP further includes one or more master machines that initiate memory access requests on the system fabric of the primary coherent system on behalf of the AP, one or more snoop machines that service requests snooped on the system fabric, and a CAPP directory having a precise directory having a plurality of entries each associated with a smaller data granule and a coarse directory having a plurality of entries each associated with a larger data granule.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

FIG. 1 is a high level block diagram of an exemplary data processing system in which a coherent device participates with a primary coherent system across a communication link through a proxy;

FIG. 2 is a more detailed block diagram of an exemplary embodiment of the data processing system of FIG. 1;

FIG. 3 is a more detailed block diagram of an exemplary embodiment of a processing unit in the data processing system of FIG. 2;

FIG. 4 is a time-space diagram of an exemplary operation on the system fabric of the data processing system of FIG. 2;

FIG. 5 is a more detailed block diagram of an exemplary embodiment of the coherent attached processor proxy (CAPP) in the processing unit of FIG. 3;

FIG. 6 is a high level logical flowchart of an exemplary process by which a CAPP coherently handles a memory access request received from an attached processor (AP) in accordance with one embodiment;

FIG. 7 is a high level logical flowchart of an exemplary process by which a CAPP coherently handles a snooped memory access request in accordance with one embodiment;

FIG. 8 is a first time-space diagram of an exemplary processing scenario in which an AP requests to coherently update a memory block within the primary coherent system to which it is attached;

FIG. 9 is a second time-space diagram of an exemplary processing scenario in which an AP requests to coherently update a memory block within the primary coherent system to which it is attached;

FIG. 10 is a third time-space diagram of an exemplary processing scenario in which an AP requests to coherently update a memory block within the primary coherent system to which it is attached;

FIG. 11 is a more detailed view of a hybrid CAPP directory in accordance with one embodiment;

FIG. 12 is a more detailed flowchart of an exemplary method of determining the CAPP coherence state in an embodiment having a hybrid CAPP directory;

FIG. 13 is a high level logical flowchart of an exemplary method by which a CAPP implements cache management commands of the AP in an embodiment having a hybrid CAPP directory;

FIG. 14 is a high level logical flowchart of an exemplary method by which a CAPP determines a partial response (Presp) for a snooped request in an embodiment having a hybrid CAPP directory;

FIG. 15 is a high level logical flowchart of an exemplary method by which an AP manages a hybrid CAPP directory; and

FIG. 16 is a data flow diagram of an exemplary design process.

DETAILED DESCRIPTION

With reference now to the figures and with particular reference to FIG. 1, there is illustrated a high level block diagram of an exemplary data processing system 100 in which a coherent device participates with a primary coherent system across a communication link through a proxy. As shown, data processing system 100 includes a primary coherent system 102 in which coherency of a distributed shared memory is maintained by implementation of a coherency protocol, such as the well-known MESI protocol or a variant thereof. The coherency protocol, which in various embodiments can be directory-based or snoop-based, is characterized by a bounded time frame in which a system-wide coherency response is determined for each memory access request.

As shown, the functionality of data processing system 100 can be expanded by coupling an attached processor (AP) 104 to primary coherent system 102 by a communication link 108. AP 104 may be implemented, for example, as a field programmable gate array (FPGA), application specific integrated circuit (ASIC), or other general or special-purpose processor or system. In various embodiments, AP 104 may, for example, serve as a co-processor that off-loads predetermined processing tasks from primary coherent system 102, provide low cost expansion of the general-purpose processing capabilities of data processing system 100, and/or provide an interface with a heterogeneous system external to primary coherent system 102. In some embodiments, AP 104 may serve as a memory controller for system memory, an I/O device, a bus performance monitor, a bus irritator (e.g., used for floor debug and system stress analysis). In support of these and other possible functions of AP 104, AP 104 preferably includes a cache 106 that holds local copies of memory blocks in the coherent memory address space of primary coherent system 102 to enable low latency access to those memory blocks by AP 104.

In many cases, the technology utilized to implement AP 104, cache 106, and/or communication link 108 has insufficient speed, bandwidth and/or reliability to guarantee that AP 104 can participate in the determination of the system-wide coherency responses for memory access requests within the bounded time frame required by the coherency protocol of primary coherent system 102. Accordingly, primary coherent system 102 further includes at least one coherent attached processor proxy (CAPP) 110 that participates on behalf of an

associated AP 104 in the determination of the system-wide coherency responses for AP 104 within a timeframe that satisfies the timing requirements of the coherency protocol of primary coherent system 102. Although not required, it is preferable if CAPP 110 is programmable and can therefore be programmed to support any of multiple different implementations of AP 104.

Referring now to FIG. 2, there is depicted a more detailed block diagram of a data processing system 200 that is one of the numerous possible embodiments of data processing system 100 of FIG. 1. Data processing system 200 may be implemented, for example, with one of the IBM Power servers, a product line of International Business Machines Corporation of Armonk, N.Y.

In the depicted embodiment, data processing system 200 is a distributed shared memory multiprocessor (MP) data processing system including a plurality of processing units 202a-202m. Each of processing units 202a-202m is supported by a respective one of shared system memories 204a-204m, the contents of which may generally be accessed by any of processing units 202a-202m. Processing units 202a-202m are further coupled for communication to a system fabric 206, which may include one or more bused, switched and/or wireless communication links. The communication on system fabric 206 includes memory access requests by processing units 202 requesting coherent access to various memory blocks within various shared system memories 204a-204m.

As further shown in FIG. 2, one or more of processing units 204a-204m are further coupled to one or more communication links 210 providing expanded connectivity. For example, processing units 202a and 202m are respectively coupled to communication links 210a-210k and 210p-210v, which may be implemented, for example, with Peripheral Component Interconnect express (PCIe) local buses. As shown, communication links 210 can be utilized to support the direct or indirect coupling of input/output adapters (IOAs) such as IOAs 212a, 212p and 212v, which can be, for example, network adapters, storage device controllers, display adapters, peripheral adapters, etc. For example, IOA 212p, which is network adapter coupled to an external data network 214, is coupled to communication link 210p optionally through an I/O fabric 216p, which may comprise one or more switches and/or bridges. In a similar manner, IOA 212v, which is a storage device controller that controls storage device 218, is coupled to communication link 210v optionally through an I/O fabric 216v. As discussed with reference to FIG. 1, communication links 210 can also be utilized to support the attachment of one or more APs 104, either directly to a processing unit 202, as is the case for AP 104k, which is coupled to processing unit 202a by communication link 210k, or indirectly to a processing unit 202 through an intermediate I/O fabric 216, as can be the case for AP 104w, which can be coupled to processing unit 202m through communication link 210v and optional I/O fabric 216v.

Data processing system 200 further includes a service processor 220 that manages the boot process of data processing system 200 and thereafter monitors and reports on the performance of and error conditions detected in data processing system 200. Service processor 220 is coupled to system fabric 206 and is supported by a local memory 222, which may include volatile (e.g., dynamic random access memory (DRAM)) and non-volatile memory (e.g., non-volatile random access memory (NVRAM) or static random access memory (SRAM)). Service processor 220 is further coupled to a mailbox interface 224 through which service processor 220 communicates I/O operations with communication link 210a.

Those of ordinary skill in the art will appreciate that the architecture and components of a data processing system can vary between embodiments. For example, other devices and

5

interconnects may alternatively or additionally be used. Accordingly, the exemplary data processing system **200** given in FIG. **2** is not meant to imply architectural limitations with respect to the claimed invention.

With reference now to FIG. **3**, there is illustrated a more detailed block diagram of an exemplary embodiment of a processing unit **202** in data processing system **200** of FIG. **2**. In the depicted embodiment, each processing unit **202** is preferably realized as a single integrated circuit chip having a substrate in which semiconductor circuitry is fabricated as is known in the art.

Each processing unit **202** includes multiple processor cores **302a-302n** for independently processing instructions and data. Each processor core **302** includes at least an instruction sequencing unit (ISU) **304** for fetching and ordering instructions for execution and one or more execution units **306** for executing instructions. The instructions executed by execution units **306** may include, for example, fixed and floating point arithmetic instructions, logical instructions, and instructions that request read and write access to a memory block in the coherent address space of data processing system **200**.

The operation of each processor core **302a-302n** is supported by a multi-level volatile memory hierarchy having at its lowest level one or more shared system memories **204** (only one of which is shown in FIG. **3**) and, at its upper levels, one or more levels of cache memory. As depicted, processing unit **202** includes an integrated memory controller (IMC) **324** that controls read and write access to an associated system memory **204** in response to requests received from processor cores **302a-302n** and operations received on system fabric **206**.

In the illustrative embodiment, the cache memory hierarchy of processing unit **202** includes a store-through level one (L1) cache **308** within each processor core **302a-302n** and a store-in level two (L2) cache **310**. As shown, L2 cache **310** includes an L2 array and directory **314**, masters **312** and snoopers **316**. Masters **312** initiate transactions on system fabric **206** and access L2 array and directory **314** in response to memory access (and other) requests received from the associated processor cores **302**. Snoopers **316** detect operations on system fabric **206**, provide appropriate responses, and perform any accesses to L2 array and directory **314** required by the operations. Although the illustrated cache

6

hierarchy includes only two levels of cache, those skilled in the art will appreciate that alternative embodiments may include additional levels (L3, L4, etc.) of private or shared, on-chip or off-chip, in-line or lookaside cache, which may be fully inclusive, partially inclusive, or non-inclusive of the contents the upper levels of cache.

As further shown in FIG. **3**, processing unit **202** includes integrated interconnect logic **320** by which processing unit **202** is coupled to system fabric **206**, as well as an instance of response logic **322**, which in embodiments employing snoop-based coherency, implements a portion of a distributed coherency messaging mechanism that maintains coherency of the cache hierarchies of processing unit **202**. Processing unit **202** further includes one or more integrated I/O (input/output) controllers **330** (e.g., PCI host bridges (PHBs)) supporting I/O communication via one or more communication links **210**. Processing unit **202** additionally includes a CAPP **110** as previously described. As shown, CAPP **110** may optionally include a dedicated I/O controller **332** (e.g., a PHB) by which CAPP **110** supports communication over an external communication link **210k** to which an AP **104k** is also coupled. In alternative embodiments, dedicated I/O controller **332** can be omitted, and CAPP **110** can communicate with AP **104** via a shared I/O controller **330**. It should be noted that there may be multiple CAPPs **110** per processing unit **202** and/or multiple processing units **202** having one or more CAPPs **110**, with each CAPP **110** independently programmed to support the operation and functionality of the attached AP **104**.

Those skilled in the art will appreciate that data processing system **200** can include many additional or alternative components. Because such additional components are not necessary for an understanding of the present invention, they are not illustrated in FIG. **3** or discussed further herein.

Referring now to FIG. **4**, there is depicted a time-space diagram of an exemplary operation on the system fabric **206** of data processing system **200** of FIG. **2** in accordance with one embodiment of a snoop-based coherence protocol. The operation begins when a master **400** (e.g., a master **312** of an L2 cache **310**, a master within an I/O controller **330** or a master in CAPP **110**) issues a request **402** on system fabric **206**. Request **402** preferably includes at least a transaction type indicating a type of desired access and a resource identifier (e.g., real address) indicating a resource to be accessed by the request. Common types of requests preferably include those set forth below in Table I.

TABLE I

Request	Description
READ	Requests a copy of the image of a memory block for query purposes
RWITM (Read-With-Intent-To-Modify)	Requests a unique copy of the image of a memory block with the intent to update (modify) it and requires destruction of other copies, if any
BKILL (Background Kill)	Requests invalidation of all cached copies of a target memory block and cancellation of all reservations for the target memory block
DCLAIM (Data Claim)	Requests authority to promote an existing query-only copy of memory block to a unique copy with the intent to update (modify) it and requires destruction of other copies, if any
DCBZ (Data Cache Block Zero)	Requests authority to create a new unique copy of a memory block without regard to its present state and subsequently modify its contents; requires destruction of other copies, if any
CASTOUT	Copies the image of a memory block from a higher level of memory to a lower level of memory in preparation for the destruction of the higher level copy
WRITE	Requests authority to create a new unique copy of a memory block without regard to its present state and immediately copy the image of the memory block from a higher level memory to a lower level memory in preparation for the destruction of the higher level copy

Further details regarding these operations and an exemplary cache coherency protocol that facilitates efficient handling of these operations may be found in U.S. Pat. No. 7,389,388, which is incorporated by reference.

Request 402 is received by snoopers 404 distributed throughout data processing system 200, including, for example, snoopers 316 of L2 caches 310, snoopers 326 of IMCs 324, and snoopers within CAPPs 110 (see, e.g., snoop machines (SNMs) 520 of FIG. 5). In general, with some exceptions, snoopers 316 in the same L2 cache 310 as the master 312 of request 402 do not snoop request 402 (i.e., there is generally no self-snooping) because a request 402 is transmitted on system fabric 206 only if the request 402 cannot be serviced internally by a processing unit 202. Snoopers 404 that receive and process requests 402 each provide a respective partial response (Presp) 406 representing the response of at least that snoopers 404 to request 402. A snoopers 326 within an IMC 324 determines the partial response 406 to provide based, for example, upon whether the snoopers 326 is responsible for the request address and whether it has resources available to service the request. A snoopers 316 of an L2 cache 310 may determine its partial response 406 based on, for example, the availability of its L2 array and directory 314, the availability of a snoop machine instance within snoopers 316 to handle the request, and the coherence state associated with the request address in L2 array and directory 314.

The partial responses 406 of snoopers 404 are logically combined either in stages or all at once by one or more instances of response logic 322 to determine a systemwide coherence response to request 402, referred to herein as a combined response (Cresp) 410. In one preferred embodiment, which will be assumed hereinafter, the instance of response logic 322 responsible for generating combined response 410 is located in the processing unit 202 containing the master 400 that issued request 402. Response logic 322 provides combined response 410 to master 400 and snoopers 404 via system fabric 206 to indicate the response (e.g., success, failure, retry, etc.) to request 402. If combined response 410 indicates success of request 402, combined response 410 may indicate, for example, a data source for a requested memory block, a cache state in which the requested memory block is to be cached by master 400, and whether "cleanup" operations invalidating the requested memory block in one or more caches are required.

In response to receipt of combined response 410, one or more of master 400 and snoopers 404 typically perform one or more actions in order to service request 402. These actions may include supplying data to master 400, invalidating or otherwise updating the coherence state of data cached in one or more caches, performing castout operations, writing back data to a system memory 204, etc. If required by request 402, a requested or target memory block may be transmitted to or from master 400 before or after the generation of combined response 410 by response logic 322.

In the following description, the partial response 406 of a snoopers 404 to a request 402 and the actions performed by the snoopers 404 in response to the request 402 and/or its combined response 410 will be described with reference to whether that snoopers 404 is a Highest Point of Coherency (HPC), a Lowest Point of Coherency (LPC), or neither with respect to the request address specified by the request. An LPC is defined herein as a memory device or I/O device that serves as the repository for a memory block. In the absence of a HPC for the memory block, the LPC holds the true image of the memory block and has authority to grant or deny requests to generate an additional cached copy of the memory block. For a typical request in the data processing system embodiment of

FIG. 2, the LPC will be the memory controller 324 for the system memory 204 holding the referenced memory block. An HPC is defined herein as a uniquely identified device that caches a true image of the memory block (which may or may not be consistent with the corresponding memory block at the LPC) and has the authority to grant or deny a request to modify the memory block. Descriptively, the HPC may also provide a copy of the memory block to a requestor in response to an operation that does not modify the memory block. Thus, for a typical request in the data processing system embodiment of FIG. 2, the HPC, if any, will be an L2 cache 310 or CAPP 110. Although other indicators may be utilized to designate an HPC for a memory block, a preferred embodiment of the present invention designates the HPC, if any, for a memory block utilizing selected cache coherency state(s), which may be held, for example, in a cache directory.

Still referring to FIG. 4, the HPC, if any, for a memory block referenced in a request 402, or in the absence of an HPC, the LPC of the memory block, preferably has the responsibility of protecting the transfer of ownership of a memory block, if necessary, in response to a request 402. In the exemplary scenario shown in FIG. 4, a snoopers 404n at the HPC (or in the absence of an HPC, the LPC) for the memory block specified by the request address of request 402 protects the transfer of ownership of the requested memory block to master 400 during a protection window 412a that extends from the time that snoopers 404n determines its partial response 406 until snoopers 404n receives combined response 410 and during a subsequent window extension 412b extending (preferably, for a programmable time) beyond receipt by snoopers 404n of combined response 410. During protection window 412a and window extension 412b, snoopers 404n protects the transfer of ownership by providing partial responses 406 to other requests specifying the same request address that prevent other masters from obtaining ownership (e.g., a retry partial response) until ownership has been successfully transferred to master 400. If necessary, master 400 may also likewise initiate a protection window 413 to protect its ownership of the memory block requested in request 402 following receipt of combined response 410.

As will be appreciated by those skilled in the art, the snoop-based coherence protocol illustrated in FIG. 4 may be implemented utilizing multiple diverse sets of coherence states. In a preferred embodiment, the cache coherence states employed within the protocol, in addition to providing (1) an indication of whether a cache is the HPC for a memory block, also indicate at least (2) whether the cached copy is unique (i.e., is the only cached copy system-wide), (3) whether and when the cache can provide a copy of the memory block to a master of a memory access request for the memory block, (4) whether the cached image of the memory block is consistent with the corresponding memory block at the LPC (system memory). These attributes can be expressed, for example, in a variant of the well-known MESI (Modified, Exclusive, Shared, Invalid) protocol including at least the coherence states summarized below in Table II.

TABLE II

Coherence state	HPC?	Unique?	Data Source?	Consistent with LPC?
M (Modified)	Yes	Yes	Yes (before Cresp)	No
T (Shared-Owner)	Yes	Unknown	Yes (after Cresp)	No
S (Shared)	No	Unknown	No	Unknown
I (Invalid)	No	No	No	N/a - data is invalid

In addition to the coherence states listed in Table II, the coherence protocol may include one or more additional transitional coherence states that can be employed, among other things, to implement protection window **412a**, window extension **412b**, and protection window **413**. For example, the coherence protocol may include an HPC Protect state that master **400** may assume in response to combined response **410** to protect transfer of HPC status (i.e., coherence ownership) to that master **400** during protection window **413**. Similarly, the coherence protocol may additionally include a Shared Protect state that a master **400** or a snoopers **404n** may assume in response to issuing or snooping a DClaim request, respectively, in order to implement protection window **413** or protection window **412a** and window extension **412b**. Further, the coherence protocol may include an Shared Protect Noted state that may be assumed to facilitate assumption of HPC status by another master **400**, as described further herein.

Referring now to FIG. 5, there is depicted a more detailed block diagram of an exemplary embodiment of the coherent attached processor proxy (CAPP) **110** in processing unit **202** of FIG. 3. As shown, CAPP **110** is coupled to interconnect logic **320** to permit CAPP **110** to transmit and receive address, control and coherency communication via system fabric **206** on behalf of (i.e., as a proxy for) an AP **104** (e.g., AP **104k**) to which it is coupled by a communication link (e.g., communication link **210k**).

CAPP **110** includes snoopers logic **500**, master logic **502**, transport logic **504**, and as discussed above, an optional I/O controller **332**. Transport logic **504** has two interfaces, a first by which transport logic **504** manages communication over communication link **210k** as necessary to comport with the messaging protocol employed by communication link **210k** and/or AP **104**, and a second by which transport logic **504** manages data communication with system fabric **206**. Thus, transport logic **504** may packetize data, may apply message encapsulation/decapsulation or encryption/decryption, may compute, append and/or verify checksums, etc., as is known in the art.

Snoopers logic **500** includes a decoder **510**, a directory **512** of the contents of the data array **552** of the cache **106** of the associated AP **104**, a snoop table **514**, a dispatcher **516**, and a set of snoop machines (SNMs) **520**. Decoder **510** of snoopers logic **500** receives memory access requests from system fabric **206** via interconnect logic **320** and optionally but preferably decodes the snooped memory access requests into a corresponding set of internal snoop requests. The set of internal snoop requests implemented by decoder **510** is preferably programmable (and in some embodiments dynamically reprogrammable) to decouple the design of CAPP **110** from that of AP **104** and to allow flexibility in mapping the memory access requests of the primary coherent system **102** to the request set of the associated AP **104**. Following decoding by decoder **510**, the target address specified by the memory access request is utilized to access directory **512** in order to look up the coherence state of the target address with respect to AP **104**. It should be noted that the coherence state indicated by directory **512** may not match or correspond to that indicated by directory **550** of cache **106** in AP **104**. Nevertheless, the use of the coherence state information in directory **512** in CAPP **110** rather than directory **550** in AP **104** enables the bounded time frame in which a system-wide coherency response is to be determined for each memory access request in primary coherent system **102** to be met, regardless of whether communication link **210** and/or AP **104** have lower speed or reliability than other components of data processing system (e.g., CAPP **110**).

The coherence state specified by directory **512** and the internal request determined by decoder **510** are then utilized by snoop table **514** to determine an appropriate partial response (Presp) to the snooped memory access request. In response to at least the internal snoop request determined by decoder **510**, coherence state output by directory **512** and Presp output by snoop table **514**, dispatcher **516** determines whether or not any further action is or may possibly be required in response to the memory access request (e.g., update of directory **512**, sourcing the target cache line to the requester, etc.), and if so, dispatches a snoop machine **520** to manage performance of that action.

Master logic **502** optionally but preferably includes a master table **530** that maps memory access and other requests originated by AP **104k** and received by CAPP **110** to internal master requests. As with the mapping performed by decoder **510** of snoopers logic **500**, the mapping performed by master table **530** decouples the design of CAPP **110** and AP **104** and enables CAPP **110** to programmably support a wide variety of diverse APs **104**. In at least some embodiments, master table **530** supports dynamic reprogramming. Master logic **502** further includes a set of master machines (MMs) **532** that services internal master requests output by master table **530**. In a typical case, a master machine **532** allocated to service an internal master request determines and manages an action to be performed to service the internal request (e.g., initiating a directory update and/or memory access request on system fabric **206**) based at least in part on the coherence state indicated for the target address of the master request by directory **512**. Data transfers to and from AP **104** via CAPP **110** in response to the operation of snoopers logic **500** and master logic **502** are tracked via operation tags allocated from tag pool **540**.

As further indicated in FIG. 5, master logic **502** includes a combined response (Cresp) table **534**. In response to receipt of a combined response representing the systemwide coherence response to a request, Cresp table **534** translates the combined response received from system fabric **206** into an internal Cresp message and distributes the internal Cresp message to master machines **532** and snoop machines **520**. Again, the translation of combined responses to internal Cresp messages by Cresp table **534** decouples the design of AP **104** from that of primary coherent system **102** and enables the interface provided by CAPP **110** to be programmable and thus support a variety of diverse APs **104**.

As noted above, several data structures (e.g., decoder **510**, snoop table **514**, master table **530** and Cresp table **534**) within CAPP **110** are preferably programmable, and in some embodiments, dynamically programmable. In one implementation, a control processor (e.g., service processor **220** or any of processing units **202** running supervisory code (e.g., hypervisor)) dynamically updates the data structures by first instructing AP **104** to invalidate its directory **550** and quiesce. The control processor then updates one or more of the data structures within CAPP **110**. In response to completion of the updates, the control processor instructs AP **104** to resume normal processing. It should also be noted that the configurations of master table **530** and snoop table **514** affects not only the mapping (translation) of incoming AP requests and snooped requests, respectively, but also the behavior of MMs **532** and SNMs **520**. That is, the behavior of MMs **532** in response to AP requests and the messages transmitted on system fabric **206** and to AP **104** are also preferably determined by the configuration of master table **530**. Similarly, the behavior of SNMs **520** in response to snooped requests and the messages transmitted on system fabric **206** and to AP **104** are preferably determined by the configuration of snoop table

11

514. Thus, the behaviors and messages of MMs 532 and SNMs 520 can be selectively changed by appropriate reprogramming of master table 530 and snoop table 514.

Referring now to FIG. 6, there is depicted a high level logical flowchart of an exemplary process by which a CAPP 110 coherently handles a memory access request received from an AP 104 in accordance with one embodiment. As with the other logical flowcharts presented herein, it should be appreciated that steps are presented in a logical rather than strictly chronological order and at least some of the illustrated steps may be performed concurrently or in a different order than that illustrated.

The process shown in FIG. 6 begins at block 600 and then proceeds to block 602, which illustrates an AP 104 generating a target address within the coherent address space of primary coherent system 102. The target address identifies a coherent storage location to which some type of access is desired, for example, an access to obtain a query-only copy of a cache line, update or invalidate contents of a storage location identified by the target address, writeback a cache line to system memory 204, invalidate a page table entry utilized to perform address translation, etc. AP 104 additionally performs a lookup of the coherence state of the target address in AP directory 550 (block 604). AP 104 then transmits to CAPP 110 a memory access request specifying the desired access, together with the coherence state read from AP directory 550 and any associated data (block 606).

The coherence state transmitted with the AP memory access request is referred to herein as the “expected state,” in that in many cases, the type of memory access request selected by AP 104 is predicated on the coherence state indicated by AP directory 550. In a preferred embodiment, AP 104 transmits the memory access request to CAPP 110 even in cases in which the expected state is or corresponds to an HPC state that, if held in an L2 cache 310, would permit the associated processor core 302 to unilaterally access the storage location identified by the target address prior to receipt of a combined response. This is the case because the coherence state determination made by AP 104 is only preliminary, with the final coherence state determination being made by CAPP 110 as described below.

In response to receipt of the AP memory access request, master table 530 of master logic 502 optionally translates the AP memory access request into an internal master request (e.g., one of the set of requests within the communication protocol specified for system fabric 206 (block 610). In a typical embodiment, the translation includes mapping the transaction type (ttype) indicated by the AP memory access request to a ttype utilized on system fabric 206. In addition, CAPP 110 determines a coherence state for the target address specified by the memory access request with respect to AP 104 (block 616). In a preferred embodiment, the coherence state is determined from multiple sources of coherence information according to a predetermined prioritization of the sources, which include (in order of increasing priority): directory 512, MMs 532 and SNMs 520. Thus, if CAPP 110 determines at block 616 that one of SNMs 520 is processing a snooped memory access request that collides with the target address, the coherence state indicated by that SNM 520 is determinative. Similarly, if CAPP 110 determines at block 616 that no SNMs 520 is active processing a request that collides with the target address, but the target address of the AP memory access request collides with the target address of a master request being processed by one of MMs 532, the coherence state indicated by that MM 532 is determinative. If

12

the request address does not collide with an active SNM 520 or MM 532, the coherence state indicated by CAPP directory 512 is determinative.

At block 620, master logic 502 determines whether or not the expected state communicated with the AP memory access request matches the coherence state determined by CAPP 110 at block 616. If so, master logic 502 allocates an MM 532 to service the AP memory access request in an Active state in which the MM 532 begins its activities to service the AP memory access request (block 621). At block 622, the MM 532 allocated to service the AP memory access request determines whether or not servicing the AP memory access request includes initiating a memory access request on system fabric 206. If not, the process passes through page connector B to block 650, which is described further below.

If, however, MM 532 determines at block 622 that servicing the AP memory access request includes initiating a memory access request on system fabric 206, the MM 532 initiates the required memory access request on system fabric 206 on behalf of AP 104 (block 624). Within a bounded time, master logic 502 receives the combined response (Cresp) for the request (block 626), which Cresp table 534 optionally translates to an internal Cresp message (block 628) and distributes to the MM 532 that initiated the memory access request. As indicated at block 630, if the combined response indicates Retry, meaning that at least one necessary participant could not service the request (e.g., was not available to service the request or was already processing another request having an address collision with the target address), the process returns to block 616, which has been described. Again determining the coherence state of CAPP 110 with respect to the target address at block 616 after a Retry combined response to the memory access request allows MM 532 to detect and respond to any changes to the coherence state occurred due to dispatch of a SNM 520 to service a conflicting request for the target address while the MM 532 was working on the memory access request. For example, if the dispatch of a SNM 520 leads to a coherence state mismatch, MM 532 may fail the requested memory access request, as discussed below with reference to block 644, or may defer servicing the memory access request and enter a Parked state, as discussed below with reference to block 642. If, on the other hand, the combined response received at block 630 indicates that the request succeeded, the MM 532 that initiated request performs any data handling actions, cleanup actions, and/or directory update actions required to complete servicing the request (block 632). The data handling actions can include, for example, MM 532 receiving requested data and forwarding the data to AP 104 or transmitting data from AP 104 on system fabric 206. The cleanup actions can include, for example, MM 532 issuing one or more kill requests on system fabric 206 to invalidate one or more copies of a cache line identified by the target address cached elsewhere within data processing system 200. The directory update actions include making any coherence update required by the request to both CAPP directory 512 and AP directory 550. Thereafter, the process shown in FIG. 6 ends at block 634.

Returning to block 620, in response to a determination that the expected coherence state specified with the AP memory access request does not match the coherence state determined by CAPP 110, the process proceeds to blocks 640-644. In one embodiment in which optional blocks 640-642 are omitted, the MM 532 allocated to the service the request transmits a Failure message to AP 104. In addition to the Failure message, MM 532 optionally further indicates, with the Failure message or in a separate directory update message, the coherence state for the target address determined by CAPP 110,

13

thus enabling AP 104 to update its AP directory 550 and to subsequently initiate an appropriate AP memory access request together with the appropriate expected state. Thereafter, the process shown in FIG. 6 ends at block 634. In this embodiment, AP 104 may require numerous requests to access the target memory block if the target memory block is highly contended by snoopers in primary coherent system 102. Accordingly, in an alternative embodiment including blocks 640-642, master logic 502 is able to increase its priority for the target memory block with respect to snoopers in primary coherent system 102 by entering a Parked state. In particular, master logic 502 determines at block 640 whether or not the coherence state mismatch detected at block 620 is due to one of SNMs 520 being active servicing a snooped memory access request that has an address collision with the target address. If not, the process proceeds to block 644, which has been described.

If, however, master logic 502 determines at block 640 that the coherence state mismatch detected at block 620 is due to one of SNMs 520 being active servicing a snooped memory access request that has an address collision with the target address, the process passes to block 642. Block 642 depicts master logic 502 allocating an MM 532 in Parked state. In the Parked state, MM 532 does not actively begin to service the AP memory access request and does not inhibit the SNM 520 that is active on the target address from completing its processing of the snooped memory access request, but does (in one embodiment) inhibit any other of the SNMs 520 and MMs 532 in the same CAPP 110 from transitioning to an active state to service a request specifying an address that collides with the target address of the AP memory access request. The allocated MM 532 remains in the Parked state until the SNM 520 that is active servicing the conflicting snooped memory access request transitions to an Idle state, and in response to this transition, itself transitions from the Parked state to an Active state. The process then passes to block 616 and following blocks, which have been described. Returning to block 616 ensures that the SNM 520 that was active on the target address did not change the CAPP coherence state from the expected state. (It should be noted, however, that the active SNM 520 may update the coherence state in CAPP directory 512 prior to its retirement, and the retirement of the SNM 520 itself will likely change the composite CAPP coherence state determined at block 616.)

In at least some embodiments, the allocation of an MM 532 in the Parked state does not absolutely inhibit any other of the SNMs 520 and MMs 532 in the same CAPP 110 from transitioning to an active state. Instead, the effects of a MM 532 in the Parked state (and/or an active state) on the dispatch of other SNMs 520 and MMs 532 to service selected types of conflicting requests can be varied, for example, via program control (i.e., via execution of an appropriate CAPP control instruction by one of processor cores 302 or AP 104) of the composite coherence state determination described above with reference to block 616. For example, to eliminate unnecessary traffic on system fabric 206, dispatcher 516 can be permitted by programmable control to dispatch a SNM 520 in an active state to service a snooped BKill request that invalidates the target memory block of a conflicting request being handled by a MM 532 in the Parked state or an active state. In cases in which another machine is dispatched to service a conflicting request while a MM 532 is in the Parked state, the MM 532 in the Parked state re-enters the Parked state when the process of FIG. 6 proceeds along the path from block 642 to blocks 616, 620, 640 and returns to block 642. Master logic 502 further preferably implements a counter to bound the number of times a MM 532 is forced to re-enter the Parked

14

state in this manner for a single AP request. When a threshold value of the counter is reached, the dispatch of other SNMs 520 and MMs 532 to service conflicting requests is then inhibited to permit the MM 532 to exit the Parked state and manage servicing of its AP request.

Referring now to block 650, in response to determining the servicing the AP memory access request does not require issuing a memory access request on system fabric 206, MM 532 updates the CAPP directory 512 as indicated by the AP memory access request. MM 532 then transmits a Success message to AP 104 to confirm the update to CAPP directory 512. The process thereafter terminates at block 632.

With reference now to FIG. 7, there is illustrated a high level logical flowchart of an exemplary process by which a CAPP 110 coherently handles a snooped memory access request in accordance with one embodiment. The illustrated process begins at block 700 and then proceeds to block 702, which depicts snooper logic 500 of CAPP 110 receiving a memory access request on system fabric 206 via interconnect logic 320. At block 704, decoder 510 decodes the snooped memory access request to determine the type of the request. In addition, at block 706, CAPP 110 determines a coherence state for the address referenced by the snooped memory access request, for example, utilizing the methodology previously described with reference to block 616.

Based on the decoded type of the snooped memory access request as determined at block 704 and the coherence state for the referenced address as determined at block 706, snooper logic 500 determines and transmits on system fabric 206 a partial response representing the coherence response of AP 104 to the snooped memory access request (block 710).

Referring now to block 712, dispatcher 516 of snooper logic 500 determines based on the partial response determined at block 710 and the decoded memory access request whether or not further action by CAPP 110 may be required to service the snooped memory access request. In general, if the coherence state determined at block 706 is Invalid, meaning that AP cache 106 does not hold a valid copy of the memory block identified by the referenced memory address, no further action on the part of CAPP 110 or AP 104 is required to service the snooped memory access request. If the coherence state determined at block 706 is other than Invalid, at least some additional action may be required on the part of CAPP 110 and/or AP 104 to service the snooped memory access request.

In response to a negative determination at block 712, the process depicted in FIG. 7 ends at block 730. If, however, dispatcher 516 determines at block 712 that further action by CAPP 110 and/or AP 104 may be required to service the snooped memory access request, dispatcher 516 dispatches one of SNMs 520 to manage any action required to service the snooped memory access request (block 714). At block 716, the dispatched SNM 520 determines whether the action required to service the snooped memory access request can be determined without the combined response representing the systemwide coherence response to the memory access request or whether the combined response is required to determine the action required to appropriately service the snooped memory access request. In response to a determination at block 716 that the combined response is not required to determine the action to perform to service the snooped memory access request, the dispatched SNM 520 manages performance of any data handling and/or directory update actions required by the decoded memory access request and coherence state to service the memory access request (block 718). Thereafter, the process illustrated in FIG. 7 ends at block 730.

15

In response to a determination at block 716 that the combined response is required to determine the action to be performed to service the snooped memory access request, the dispatched SNM 520 waits for the combined response, as shown at block 720. In response to receiving the combined response, Cresp table 534 optionally translates the combined response into an internal Cresp message employed by CAPP 110 (block 722). The dispatched SNM 520 then manages performance of any data handling and/or directory update actions required by the combined response to service the memory access request (block 724). Thereafter, the process illustrated in FIG. 7 ends at block 730.

Referring now to FIG. 8, there is depicted a first time-space diagram of an exemplary processing scenario in which an AP 104 requests to coherently update a memory block within the primary coherent system 102 to which it is attached. For purposes of illustration, the exemplary processing scenario given in FIG. 8 and other similar figures will be described with reference to the illustrative hardware embodiments given in FIGS. 2-3 and 5.

As the exemplary processing scenario begins, an AP 104 processes a command (e.g., a software or firmware instruction executed within AP 104) specifying an update to a memory block identified by a target address within the coherent address space of primary coherent system 102. In response to the command, AP 104 allocates one of its idle finite state machines (FSMs) to manage performance of the command and performs a lookup of the target address in AP directory 550, as indicated by arrow 800. The AP FSM transitions from an idle state (indicated by "X") to an Update Active state and, based on a determination that the target address has an Invalid coherence state with respect to AP directory 550, transmits to CAPP 110 an update request with an expected state of Invalid, as shown at reference numeral 802.

In response to receipt from AP 104 of the update request, CAPP 110 translates the AP update request into a RWITM request, which as indicated in Table I, is one of the set of requests within the communication protocol specified for system fabric 206. In addition, CAPP 110 determines a coherence state for the target address specified by the memory access request. Because in this case, the target address of the RWITM request does not collide with an address that an MM 532 or SNM 520 is currently processing, the coherence state of the target address for CAPP 110 is determined by CAPP directory 512, which returns Invalid.

The previously idle MM 532 allocated to service the RWITM request, in response to determining a coherence state match between the expected state and the coherence state determined by CAPP 110, transitions to a Valid state and initiates the RWITM request on system fabric 206 as shown at reference numeral 806. The RWITM request requests a copy of the target memory block and further requests invalidation of all other cached copies of the memory block (to permit AP 104 to modify the memory block). Within a bounded time, MM 532 receives a combined response indicating success of the RWITM request, as indicated at reference numeral 808. MM 532 also receives a copy of the requested memory block, possibly prior to, concurrently with, or after the combined response.

In response to receiving the combined response indicating success of the RWITM request, MM 532 transitions to the HPC Protect state, thus initiating a protection window 413 for the target address. In addition, as indicated by arrow 810, MM 532 updates the coherence state for the target address in CAPP directory 512 to Modified. In addition, as indicated by arrow 812, MM 532 transmits the copy of the requested

16

memory block and a Complete message to AP 104. Thereafter, MM 532 returns to the Idle state. In response to receipt of the requested memory block and Complete message, the AP FSM directs the requested update to the target memory block, storage of the updated target memory block in array 552, and update of the coherence state for the target address in AP directory 550 to Modified. The updates to AP cache 106 are performed asynchronously to the update to CAPP directory 512, and due to the possibly unreliable connection provided by communication link 210, may require CAPP 110 to retransmit the Complete message one or more times. Thereafter, the AP FSM returns to the Idle state.

It can also be appreciated by reference to FIG. 8 that (depending on the presence or absence of other colliding requests) the processing of a read request of AP 104 could be handled similarly to the illustrated processing scenario, with the following exceptions: the AP FSM would assume the Read Active state rather than the Update Active state, MM 532 would assume the Shared Protect state following receipt of the combined response indicated by arrow 808 rather than the HPC Protect state, and CAPP directory 512 and AP directory 550 would be updated to the Shared state rather than the Modified State.

With reference now to FIG. 9, there is depicted a second time-space diagram of an exemplary processing scenario in which an AP 104 requests to coherently update a memory block within the primary coherent system 102 to which it is attached.

As the exemplary processing scenario begins, an AP 104 processes a command (e.g., a software or firmware instruction executed within AP 104) specifying an update to a memory block identified by a target address within the coherent address space of primary coherent system 102. In response to the command, AP 104 allocates one of its idle finite state machines (FSMs) to manage performance of the command and performs a lookup of the target address in AP directory 550, as indicated by arrow 900. The AP FSM transitions from an Idle state (indicated by "X") to an Update Active state and, based on a determination that the target address has an Shared-Owner (T) coherence state with respect to AP directory 550, transmits to CAPP 110 an update request with an expected state of T, as shown at reference numeral 902.

In response to receipt from AP 104 of the update request, CAPP 110 translates the update request to a BKill request. As described above with reference to Table I, the BKill request requests invalidation of all other cached copies of the memory block to permit AP 104 to modify its existing HPC copy of the target memory block. CAPP 110 additionally determines a coherence state for the target address specified by the update request with respect to CAPP 110, as shown at reference numeral 904. Because in this case, the target address of the update request collides with an address that a SNM 520 is currently processing, the state of that SNM 520 is determinative, meaning that CAPP 110 determines an HPC Protect state. Thus, the coherence state determined by CAPP 110 does not match the expected state. In embodiments in which the optional functionality described above with reference to blocks 640-642 of FIG. 6 is not implemented, CAPP 110 would respond to the update request by transmitting a Failure message to AP 104. However, in the illustrated case in which the optional functionality described above with reference to blocks 640-642 of FIG. 6 is implemented, CAPP 110 allocates an idle MM 532 to service the BKill request in the Parked state, as indicated by arrow 906. As noted above, the Parked state of the MM 532 inhibits any other SNM 520 from

17

transitioning to an active state to service a snooped memory access request for the target address.

In response to the SNM 520 that is active working on the conflicting address transitioning to the Idle state without modifying the matching T coherence state in CAPP directory 512 (e.g., as would be the case if the snooped memory access request is a Read request), the MM 532 verifies that the coherence state determined for CAPP 110 (which is the T state recorded in CAPP directory 512 in the absence of a SNM 520 or MM 532 active on a conflicting address) matches the expected state, as discussed previously with reference to block 616 of FIG. 6. In response to verifying that the coherence state of CAPP directory 110 matches the expected state, the MM 532 allocated to service the BKill request transitions to the HPC Protect state (thus initiating a protection window 413 for the target address) and initiates the BKill request on system fabric 206 as shown at reference numeral 910. In other scenarios (not illustrated) in which SNM 520 modifies the coherence state in CAPP directory 512 (e.g., as would be the case if the snooped memory access request is a RWITM request), MM 532 instead returns a failure message to AP 104 and returns to the Idle state.

Returning to the scenario shown in FIG. 9, in response to the BKill request, MM 532 receives a combined response indicating success of the BKill request, as indicated at reference numeral 912. In response to receiving the combined response indicating success of the BKill request, MM 532 updates the coherence state for the target address in CAPP directory 512 to Modified. In addition, as indicated by arrow 914, MM 532 transmits a Complete message to AP 104. Thereafter, MM 532 returns to the Idle state. In response to receipt of the Complete message, the AP FSM directs the update of the coherence state for the target address in AP directory 550 from T to Modified and the update of the corresponding cache line in AP array 552. Thereafter, the AP FSM returns to the Idle state.

Referring now to FIG. 10, there is depicted a third time-space diagram of an exemplary processing scenario in which an AP 104 requests to coherently update a memory block within the primary coherent system 102 to which it is attached.

As the exemplary processing scenario shown in FIG. 10 begins, an AP 104 processes a command (e.g., a software or firmware instruction executed within AP 104) specifying an update to a memory block identified by a target address within the coherent address space of primary coherent system 102. In response to the command, AP 104 allocates one of its idle finite state machines (FSMs) to manage performance of the command and performs a lookup of the target address in AP directory 550, as indicated by arrow 1000. The AP FSM transitions from an Idle state (indicated by "X") to an Update Active state and, based on a determination that the target address has an Shared (S) coherence state with respect to AP directory 550, transmits to CAPP 110 an update request with an expected state of S, as shown at reference numeral 1002.

In response to receipt from AP 104 of the update request, CAPP 110 translates the update request to a DClaim request. As described above with reference to Table I, the DClaim request requests invalidation of all other cached copies of the target memory block to permit AP 104 to modify its existing Shared copy of the target memory block. CAPP 110 additionally determines a coherence state for the target address specified by the update request with respect to CAPP 110, as shown at reference numeral 1004. Because in this case, the target address of the update request collides with an address of a snooped DClaim request that a SNM 520 is currently processing, the state of that SNM 520 is determinative, meaning

18

that CAPP 110 determines the Shared Protect (SP) state. Thus, the coherence state determined by CAPP 110 does not match the expected state of Shared (see, e.g., block 620 of FIG. 6). Consequently, CAPP 110 allocates an idle MM 532 to the DClaim request in the Parked (P) state, as illustrated by arrow 1006 and as previously described with reference to block 642 of FIG. 6.

In response to the snooped DClaim request, the SNM 520 that is active working on the snooped DClaim request updates the coherence state of the target address in CAPP directory 512 to the Shared Protect Noted state, as indicated by arrow 1010, and additionally transmits a Kill message to AP 104 to cause the coherence state in AP directory 550 to be updated to the Invalid state, as indicated by arrow 1012. As shown in FIG. 10, the SNM 520 thereafter returns to the Idle state.

In response to the SNM 520 returning to the Invalid state, the MM 532 allocated to the DClaim request transitions from the Parked state to an active state and again determines the coherence state of the target memory address with respect to CAPP 110, as described above with reference to block 616 of FIG. 6. Because the Parked state inhibits the dispatch of any other SNM 520 to service a conflicting address, the coherence state specified by CAPP directory 512 (i.e., Shared Protect Noted) is determinative of the coherence state of the target memory address with respect to CAPP 110. In response to detecting a mismatch of the coherence state in CAPP directory 512 (Shared Protect Noted) with the expected state (Shared), the MM 532 provides a Failure message to AP 104 to indicate failure of the DClaim request of AP 104, as indicated by arrow 1014.

Due to the potential unreliability of communication link 210, the invalidation in AP directory 550 initiated by SNM 520 is preferably confirmed by receipt of MM 532 of a full handshake from AP 104 as indicated by arrow 1018. If MM 532 does not receive a handshake from AP 104 confirming invalidation of the target memory address in AP directory 550 within a predetermined time period, MM 532 preferably retries a Kill message until the handshake is returned by AP 104 or a failure threshold is reached. In response to receipt of the handshake from AP 104, the MM 532 allocated to the DClaim request returns to the Idle state.

As will be appreciated, in an alternative embodiment, CAPP 110 can instead accommodate for the possible unreliability of communication link 210 by leaving the SNM 520 allocated to service the conflicting DClaim request in the Shared Protect state until the SNM 520 receives the handshake from AP 104. However, this alternative embodiment consumes more resources in that it requires both the SNM 520 and MM 532 to remain active for longer periods of time, thus reducing the availability of resources to service other memory access requests received from AP 104 or snooped on system fabric 206.

The AP FSM, in response to receiving Kill message 1012, transitions from the Update Active state to a Kill Active state, reflecting a need to invalidate the target memory block in CAPP directory 512. Accordingly, the AP FSM performs a lookup in AP directory 550 (as indicated by arrow 1020) transmits a Kill request 1022 to CAPP 110 specifying the same target memory address as its earlier update request and indicating an expected coherence state of Shared Protect Noted (which the AP FSM received in Kill message 1012). In response to the Kill request, master logic 502 again determines the coherence state of the target memory address with respect to CAPP 110 as described above with respect to block 616 of FIG. 6, and as indicated in FIG. 10 by arrow 1024. In response, to determining that the coherence state of the target memory address with respect to CAPP 110 (i.e., the Shared

19

Protect Noted state indicated by CAPP directory 512) matches the expected state indicated by AP 104, master logic 502 allocates a MM 532 (which could be the same MM 532 or a different MM 532) in an Active (A) state to service the AP Kill request, as illustrated by arrow 1026 and described above with reference to block 621 of FIG. 6. Because the Kill request does not require a memory access request to be issued on system fabric 206, the MM 532 updates the CAPP directory 512 as indicated by the AP memory access request, as described above with reference to block 650 of FIG. 6, in this case by invalidating the target memory address in CAPP directory 512. This update to CAPP directory 512 is illustrated in FIG. 10 by arrow 1028. On completion of the update to CAPP directory 512, MM 532 also transmits a Success message to AP 104 to confirm the update to CAPP directory 512, as indicated in FIG. 10 by arrow 1030 and as described above with respect to block 652 of FIG. 6.

After the scenario illustrated in FIG. 10, the processing scenario illustrated in FIG. 8 can be performed in order to allow AP 104 to update the target memory block of primary coherent system 102.

One common computational model in multiprocessor data processing systems, such as data processing systems 100 and 200, is a producer-consumer model in which one or more threads of execution produce a data set and one or more other threads consume the data set. Assuming proper software synchronization of the activity of the various threads, for example, utilizing control blocks, interrupts, locks, messaging or other synchronization constructs, a thread of execution in a producer-consumer model should experience little or no contention from other threads for cachelines within the data set on which it is working. In at least some use scenarios, one or more APs 104 can be employed as producers and/or consumers of the data set to perform computational functions such as encryption or decryption, matrix transformations, textual translation, etc.

In such use scenarios, it would be advantageous if the CAPP 110 participating in coherent communication in primary coherent system 102 on behalf of the AP 104 could protect the data set or particular regions thereof from conflicting accesses (e.g., by other APs 104 or processing units 202) while the AP 104 is working on the data set. With a snooped-based coherence protocol, protecting the data set entails tracking the addresses of the data structure in CAPP directory 512 and providing Retry partial responses to conflicting accesses. However, because the data set on which an AP 104 is working can be large (even if the data structure containing the data set is sparsely populated and the data set is comparatively small), it can be impractical from a cost and performance standpoint to grow CAPP directory 512 to a sufficient size to track the entire data structure if each directory entry is utilized to track a small coherence granule (e.g., a single cacheline containing 128 bytes of data). However, it is similarly impractical from a performance standpoint for each directory entry in CAPP directory 512 simply to track large coherence granules (e.g., memory pages) because the amount of coherency communication overhead (e.g., invalidations and retries) required to obtain authority to update the coherence granules. Consequently, in at least one embodiment, CAPP directory 512 is preferably implemented with a hybrid coarse/precise structure providing multiple (e.g., two) levels of directory for the contents of a single cache array (i.e., AP array 552).

With reference now to FIG. 11, there is illustrated a more detailed view of a hybrid CAPP directory 512 in accordance with one embodiment. In the depicted embodiment, CAPP directory 512 includes a coarse directory 1100 and a precise

20

directory 1102. In the depicted example, coarse directory 1100 is a four way set-associative directory including way 0 1110a, way 1 1110b, way 2 1110c, and way 3 1110d. In one example, each of ways 1110a-1110d may have a capacity of 256 entries, each having two sectors. Of course, other embodiments may have more or fewer ways, entries, and/or sectors. When accessed by the target address of a memory access request, an index field of the target address is utilized to index into each of ways 1110a-1110d and to read out the contents the relevant entry. A series of multiplexers 1112 receives the directory entries read out from ways 1110a-1110d and selects one of the sectors of the entries based on a sector select field from the target address (which in the illustrated example of two sectors can include a single bit). A series of comparators 1114 compares the address tags recorded in the sectors selected by multiplexers 1112 with an address tag field of the target address to make a hit/miss determination for the target address in coarse directory 1100.

As further depicted in FIG. 11, the depicted embodiment of precise directory 1102 is an eight way set-associative directory including ways 0-7, each including two ranks 0-1. Thus, in this embodiment, precise directory 1102 includes sixteen directory arrays, which are identified in rank 0 by reference numerals 1122a0-1122h0 and are identified in rank 1 by reference numerals 1122a1-1122h1. In one example, each of the directory arrays may have a capacity of 256 entries. Of course, other embodiments may have more or fewer ways, entries and/or ranks. When accessed by a target address of a memory access request, an index field of the target address is utilized to index into each of directory arrays 1122a0-1122h0 and 1122a1-1122h1 to read out the contents the relevant entry. A series of multiplexers 1124 receives the directory entries read out from the directory arrays and selects the entries of one of the ranks based on a rank select field from the target address (which in the illustrated example of two ranks can include a single bit). A series of comparators 1126 compares the address tags recorded in the entries selected by multiplexers 1124 with an address tag field of the target address to make a hit/miss determination for the target address in precise directory 1102.

It should further be understood that in some embodiments, directory contention in CAPP directory 512 can be reduced by implementing multiple coarse directories 1100 and multiple precise directories 1102, for example, a coarse directory 1100 and precise directory 1102 for even request addresses and a coarse directory 1100 and precise directory 1102 for odd request addresses. Directory contention can alternatively or additionally further be reduced by implementing each coarse directory 1100 and precise directory 1102 as a dual ported directory, with separate ports for read and write accesses. Further, in some embodiments, CAPP directory 512 includes or is associated with a configuration register 1130 having a plurality of settings that can be set by hardware and/or software to determine the sizes and configurations of coarse directory 1100 and precise directory 1102. For example, in one embodiment, a precise directory 1102 having eight ways, two ranks and 256 entries per directory array can be configured by appropriate settings of configuration register 1130 to have a maximum addressable capacity of 1 MB, 512 kB, or 256 kB. Similarly, in one embodiment, a coarse directory 1100 having four ways and 256 entries per way can be configured with either (1) two 4 kB sectors per entry, which yields a 2 MB addressable capacity per way and 8 MB total addressable capacity, or (2) four 64 kB sectors per entry, which yields 64 MB addressable capacity per way and 256 MB total addressable capacity.

21

In a preferred embodiment, each CAPP directory **512** and its constituent coarse directory **1100** and precise directory **1102** is preferably independently configured and controlled by the associated AP **104**. Thus, the AP **104** determines the configurations of coarse directory **1100** and precise directory **1102** and further controls which of coarse directory **1100** and precise directory **1102** will be employed to track each cacheline of interest to AP **104** (whether or not such cachelines are held in array **552** of AP cache **106**). In various scenarios, AP **104** can specify that an entry for a cacheline is held only in coarse directory **1100** (e.g., in an entry corresponding to numerous cachelines), is held only in precise directory **1102** (e.g., in an entry corresponding to only a single cacheline), or is held in both coarse directory **1100** and precise directory **1102**. In at least some embodiments, an AP **104** can specify which of directories **1100**, **1102** is/are to track a given cacheline (and can change which directory or directories track a given cacheline) utilizing explicit special-purpose directory write commands directed to the associated CAPP **110** or in one or more fields of AP requests (e.g., read, RWITM, DClaim, etc.) issued by the AP **104** to its associated CAPP **110**.

With proper management, contention for cachelines of interest to an AP **104** that have an associated entry in coarse directory **1100** should be relatively low. However, contention for such cachelines may occasionally occur. To avoid contention for a few cachelines from unnecessarily invalidating an entry in coarse directory **1100** representing numerous cachelines, AP **104** preferably implements a directory management policy that permits entries in coarse directory **1100** to be “holey,” that is, to represent with a data-valid coherence state a block of multiple contiguous cache lines, not all of which are valid in AP cache **550**. This directory management policy also enables AP **104** to establish, in coarse directory **1100**, one or more entries corresponding to a working data set prior to AP **104** retrieving and/or producing the corresponding cachelines of data. As noted further below, entries in coarse directory **1100**, whether “holey” or not, preferably represent the corresponding collection of multiple contiguous cachelines with a single coherence state corresponding to the highest coherence state of any of the associated cachelines (e.g., M, T, S and I in descending order in the exemplary set of coherence states given in Table II).

Referring now to FIG. 12, there is depicted a more detailed flowchart of an exemplary method of determining the CAPP coherence state in an embodiment having a hybrid CAPP directory **512** including a coarse directory **1100** and a precise directory **1102**. The illustrated process can be performed, for example, at block **616** of FIG. 6 and at block **706** of FIG. 7.

The process of FIG. 12 begins at block **1200** and then proceeds to blocks **1202-1208**, which illustrate CAPP **110** determining whether or not a target address (i.e., the address specified by an AP request or a snooped request) hits one of SNMs **520**, MMs **532**, precise directory **1102** and coarse directory **1100**. If the target address does not hit any of SNMs **520**, MMs **532**, precise directory **1102** and coarse directory **1100**, then CAPP **110** determines that the CAPP coherence state for the target address is invalid (block **1210**). The process of FIG. 12 then terminates at block **1220**.

If, on the other hand, the target address hits one or more of SNMs **520**, MMs **532**, precise directory **1102** and coarse directory **1100**, CAPP **110** applies a prioritization to the results of the lookup of the target address. For example, if CAPP **110** determines at block **1202** that the target address hits one of SNMs **520**, then CAPP **110** utilizes, as the CAPP coherence state, the coherence state indicated by the SNM **520** working on the target address (block **1212**). If CAPP **110**

22

determines at block **1204** that the target address missed SNMs **520** and hit one of MMs **532**, then CAPP **110** utilizes, as the CAPP coherence state, the coherence state indicated by the MM **532** working on the target address (block **1214**). If CAPP **110** determines at block **1206** that the target address missed SNMs **520** and MMs **532** and hit an entry in precise directory **1102**, then CAPP **110** utilizes, as the CAPP coherence state, the coherence state indicated by the matching entry in precise directory **1102** (block **1216**). Finally, if CAPP **110** determines at block **1208** that the target address missed SNMs **520**, MMs **532** and precise directory **1102** and hit in coarse directory **1100**, then CAPP **110** utilizes, as the CAPP coherence state, the coherence state indicated by the matching entry in coarse directory **1100** (block **1218**). Following any of blocks **1212-1218**, the process of FIG. 12 ends at block **1220**.

With reference now to FIG. 13, there is illustrated a high level logical flowchart of an exemplary method by which a CAPP **110** implements cache management commands of the AP **104** in an embodiment having a hybrid CAPP directory **512** including a coarse directory **1100** and a precise directory **1102**. The process of FIG. 13 may be performed, for example, in response to an explicit directory write command or as part of the directory update operations performed at block **632** of FIG. 6 in response to a successful AP request.

The process of FIG. 13 begins at block **1300** and then proceeds to block **1302**, which illustrates CAPP **110** determining which directory or directories were specified by the associated AP **104** (e.g., in a directory write command or in an AP request) to track a given memory block containing one or more cachelines. If AP **104** specified that the memory block be tracked by coarse directory **1100**, then, if required (i.e., if such an entry does not already reside in coarse directory **1100**), CAPP **110** installs an entry corresponding to the memory block in coarse directory **1100**, as shown at block **1304**. For example, if AP **104** determines that it is to produce or consume a data set comprising numerous contiguous cachelines of data, AP **104** may instruct CAPP **110** to install a corresponding entry representing the contiguous cachelines in coarse directory **1100**. If, for example, AP **104** is to consume the data set, AP **104** may instruct CAPP **110** to establish the entry in coarse directory **1100** in the first of a series of RWITM requests that load the data set into AP cache **550** for processing by AP **104**. As shown at block **1306**, CAPP **110** preferably sets the coherence state field of the entry in coarse directory **1100** to a coherence state reflect the highest coherence state of all the cachelines in the associated memory block. Thus, if the coherence state of at least one cacheline in the associated memory block is Modified (as would be the case following a RWITM request successfully loading an initial cacheline of a set of multiple contiguous cachelines), the coherence state field of the entry in coarse directory **1100** would be set to the Modified coherence state at block **1306**. Following block **1306**, the process of FIG. 13 ends at block **1320**.

In response to a determination at block **1302** that AP **104** has specified that an entry is to be installed in precise directory **1102** (e.g., solely or in addition to the installation of an entry in coarse directory **1100**), CAPP **110** installs an entry corresponding to a particular cacheline into precise directory **1102**, if needed (block **1310**). AP **104** may opt to install an entry in precise directory **1102**, for example, if contention for the cacheline is experienced or if AP **104** has completed producing a data set including the cacheline and is now ready for the cacheline to be sourced via CAPP **110** to one or more consumers in data processing system **200**. As shown at block **1312**, CAPP **110** additionally sets the coherence field of the entry in precise directory **1102** as indicated, for example, by

23

a combined response for a request targeting the cacheline or as specified by AP 104. Thereafter, the process of FIG. 13 ends at block 1320.

Referring now to FIG. 14, there is depicted a high level logical flowchart of an exemplary method by which a CAPP 110 determines an partial response (Presp) for a snooped request in an embodiment having a hybrid CAPP directory 512 including a coarse directory 1100 and a precise directory 1102. The illustrated process may be performed, for example, in response to a CAPP 110 snooping a memory access request, as illustrated block 710 of FIG. 7.

The process of FIG. 14 begins at block 1400 and then proceeds to block 1402, which depicts a determination of whether or not the target address of the snooped memory access request hit or missed in SNMs 1202, MMs 1204 and CAPP directory 512, as illustrated, for example, at block 706 of FIG. 7 and at blocks 1202-1208 of FIG. 12. In response to a determination that the target address of the snooped memory access request missed in CAPP 110, CAPP 110 may provide a Null Presp, as shown at block 1404. Following block 1404, the process of FIG. 14 ends at block 1412.

Referring again to block 1402, in response to a determination at the target address of the snooped memory access hit in CAPP 110, CAPP 110 determines the appropriate Presp based on the snooped memory access request and the composite CAPP coherence state determined according to the process of FIG. 12, as depicted at block 1406. For memory access requests that hit solely in coarse directory 1100, CAPP 110 cannot intervene the target cacheline or update the coherence state of the target cacheline in coarse directory 1100 because the coarse directory entry represents a granule including multiple cachelines. Accordingly, as shown at blocks 1408-1410, in at least cases in which the memory access request hit solely in coarse directory 1100, CAPP 110 notifies AP 104 of the hit in coarse directory 1100, thus indicating to AP 104 a need to install an entry for the target cacheline in precise directory 1102. Thereafter, in response to a directory write command from AP 104 instructing CAPP 110 to install an entry for the target address in precise directory 1102, CAPP 110 installs an entry in precise directory 1102 utilizing the coherence state specified by AP 104 (AP 104 tracks the precise state of the relevant cacheline in its AP directory 550). The installation of the entry in precise directory 1102 enables CAPP 110 to thereafter service the memory access request if retried by the requester.

For example, if the target address of a snooped RWITM or DClaim request hit solely in coarse directory 1100 in a non-HPC coherence state (e.g., S), then CAPP 110 determines a Shared Presp at block 1406 (where the Shared Presp indicates that AP cache 106 may possibly hold a copy of the target cacheline) and notifies AP 104 at block 1410. In response to the Shared Presp, the requesting master initiates a background kill (BKill) operation, which continues until AP 104 installs the target cacheline into precise directory 1102. Once the target cacheline is installed in coarse directory 1102, the BKill operation is allowed to complete, either because the individual target cacheline is updated by CAPP 110 from Shared to Invalid in precise directory 1102 or because the target cacheline was not in (i.e., was a hole in) the coarse granule and was installed by AP 104 into coarse directory 1102 in the Invalid coherence state.

As another example, if the target address of a snooped memory access request hit solely in coarse directory 1100 in an HPC coherence state (e.g., M or T), then CAPP 110 determines a Retry Presp at block 1406 and notifies AP 104 at block 1410. The Retry Presp is provided in this case because CAPP 110 cannot appropriately service a memory access

24

request directed to a particular target cacheline by reference to the coherence state of the coarse directory entry of the granule containing the target cacheline, as the granule's coherence state may not be the actual coherence state of the individual target cacheline. In response to the Retry Presp, which would indicate a possibly protecting HPC to the requesting master, the requesting master continues to reissue the memory access request for the target cacheline until AP 104 installs the target cacheline into precise directory 1102. Once the target cacheline is installed into precise directory 1102, CAPP 110 can either intervene the target cacheline (because the target cacheline was held in the coarse granule), or the memory access request is serviced by the LPC (e.g., a system memory 204) because the target cacheline was not in (i.e., was a hole in) the coarse granule and was installed into coarse directory 1102 in the Invalid coherence state by AP 104.

Following block 1408 or 1410, the process of FIG. 14 ends at block 1412.

With reference now to FIG. 15, there is illustrated a high level logical flowchart of an exemplary method by which an AP 104 manages a hybrid CAPP directory 512 including a coarse directory 1100 and a precise directory 1102. The process of FIG. 15 begins at block 1500 and then proceeds to block 1502, which depicts AP 104 determining whether or not AP 104 has received notification from CAPP 110 of a snooped memory access request that hit in coarse directory 1100, for example, as described above with reference to block 1410. (It should be appreciated that CAPP 110 may supply such notifications to AP 104 in cases in which the snooped memory access request hits in one or more of SNMs 520, MMs 532 and precise directory 1102.) In response to a negative determination, the process iterates at block 1502. If, however, AP 104 determines at block 1502 that it has received a notification from the associated CAPP 110 of a snooped memory access request that hit in coarse directory 1100, AP 104 determines at block 1504 whether or not the target address of the memory access request is a hole in the coarse granule, that is, whether or not AP cache directory 550 indicates an Invalid state for the target address. In response to an affirmative determination at block 1504, AP 104 issues a directory write command directing CAPP 110 to install into precise directory 1102 a precise entry for the target cacheline indicating the Invalid coherence state (block 1506). Thereafter, the process of FIG. 15 ends at block 1520.

Returning to block 1504, in response to AP 104 determining that the target cacheline is not a hole in the coarse granule, the process proceeds to block 1508, which illustrates AP 104 issuing a directory write command directing CAPP 110 to install into precise directory 1102 a precise entry for the target cacheline using the coherence state of the coarse directory entry. In addition, AP 104 marks the target cacheline as a hole in the coarse granule, for example, in a bit vector in which each bit represents a corresponding cacheline in the coarse granule. As will be appreciated, the number of bits marked in the bit vector (e.g., set to a predetermined one of "1" or "0") provides a measure of how "holey" the granule associated with the coarse directory entry is. At block 1510, AP 104 may additionally determine whether or not the number of holes (i.e., invalid cachelines) in the coarse granule satisfies (e.g., is equal to or greater than) a decomposition threshold. This determination can be made, for example, by a population count of bits set to a selected state (i.e., "1" or "0") in the corresponding bit vector. In response to a negative determination at block 1510, the process of FIG. 15 ends at block 1520. If, however, AP 104 determines at block 1510 that the decomposition threshold is satisfied for the coarse granule

25

containing the target cacheline, AP 104 instructs CAPP 110 to decompose the corresponding entry in coarse directory 1100. In particular, AP 104 instructs CAPP 110 to install in precise directory 1102 one or more entries corresponding to valid cachelines within the associated coarse granule (e.g., an entry for each cacheline held in AP cache 106 in a data-valid coherence state) and to invalidate the relevant entry in coarse directory 1100. Following block 1512, the process of FIG. 15 ends at block 1520.

In an alternative process, at block 1512 AP 104 may instead be configured to or may choose to (e.g., based on dynamic workload or a selected operating mode) initiate writeback to system memory 204 of any modified cachelines within the granule represented by the entry in coarse directory 1100, creating a progressively more “holey” entry. In response to completion of the writeback of any modified cachelines in the coarse granule represented by the entry in coarse directory 1100, AP 104 can then invalidate the entry in coarse directory 1100, thus freeing CAPP 110 from the responsibility for intervening requested cachelines in the granule and allowing any requesters to satisfy memory access requests targeting the granule from system memory 204 (or another intervening cache).

FIG. 15 illustrates a particular implementation of demand-based directory management process in which AP 104 performs directory management operations on coarse directory 1100 in response to receipt of notification by CAPP 110 of snooped memory access requests that hit in coarse directory 1100. It should be appreciated that the described process is merely exemplary and alternative or additional directory management techniques, including those initiated by AP 104 in the absence of a snooped memory access request hitting in coarse directory 1100, may be employed. Further, it should be appreciated that AP 104 may optionally employ a corresponding process to compose coarse directory entries from multiple precise directory entries as needed to accommodate its current workload and to free storage space in precise directory 1102.

Referring now to FIG. 16, there is depicted a block diagram of an exemplary design flow 1600 used for example, in semiconductor IC logic design, simulation, test, layout, and manufacture. Design flow 1600 includes processes, machines and/or mechanisms for processing design structures or devices to generate logically or otherwise functionally equivalent representations of the design structures and/or devices described above and shown in FIGS. 1-3, 5 and 11. The design structures processed and/or generated by design flow 1600 may be encoded on machine-readable transmission or storage media to include data and/or instructions that when executed or otherwise processed on a data processing system generate a logically, structurally, mechanically, or otherwise functionally equivalent representation of hardware components, circuits, devices, or systems. Machines include, but are not limited to, any machine used in an IC design process, such as designing, manufacturing, or simulating a circuit, component, device, or system. For example, machines may include: lithography machines, machines and/or equipment for generating masks (e.g. e-beam writers), computers or equipment for simulating design structures, any apparatus used in the manufacturing or test process, or any machines for programming functionally equivalent representations of the design structures into any medium (e.g. a machine for programming a programmable gate array).

Design flow 1600 may vary depending on the type of representation being designed. For example, a design flow 1600 for building an application specific IC (ASIC) may differ from a design flow 1600 for designing a standard com-

26

ponent or from a design flow 1600 for instantiating the design into a programmable array, for example a programmable gate array (PGA) or a field programmable gate array (FPGA) offered by Altera® Inc. or Xilinx® Inc.

FIG. 16 illustrates multiple such design structures including an input design structure 1620 that is preferably processed by a design process 1610. Design structure 1620 may be a logical simulation design structure generated and processed by design process 1610 to produce a logically equivalent functional representation of a hardware device. Design structure 1620 may also or alternatively comprise data and/or program instructions that when processed by design process 1610, generate a functional representation of the physical structure of a hardware device. Whether representing functional and/or structural design features, design structure 1620 may be generated using electronic computer-aided design (ECAD) such as implemented by a core developer/designer. When encoded on a machine-readable data transmission, gate array, or storage medium, design structure 1620 may be accessed and processed by one or more hardware and/or software modules within design process 1610 to simulate or otherwise functionally represent an electronic component, circuit, electronic or logic module, apparatus, device, or system such as those shown in FIGS. 1-3, 5 and 11. As such, design structure 1620 may comprise files or other data structures including human and/or machine-readable source code, compiled structures, and computer-executable code structures that when processed by a design or simulation data processing system, functionally simulate or otherwise represent circuits or other levels of hardware logic design. Such data structures may include hardware-description language (HDL) design entities or other data structures conforming to and/or compatible with lower-level HDL design languages such as Verilog® and VHDL, and/or higher level design languages such as C or C++.

Design process 1610 preferably employs and incorporates hardware and/or software modules for synthesizing, translating, or otherwise processing a design/simulation functional equivalent of the components, circuits, devices, or logic structures shown in FIGS. 1-3, 5 and 11 to generate a netlist 1680 which may contain design structures such as design structure 1620. Netlist 1680 may comprise, for example, compiled or otherwise processed data structures representing a list of wires, discrete components, logic gates, control circuits, I/O devices, models, etc. that describes the connections to other elements and circuits in an integrated circuit design. Netlist 1680 may be synthesized using an iterative process in which netlist 1680 is resynthesized one or more times depending on design specifications and parameters for the device. As with other design structure types described herein, netlist 1680 may be recorded on a machine-readable storage medium or programmed into a programmable gate array. The medium may be a non-volatile storage medium such as a magnetic or optical disk drive, a programmable gate array, a compact flash, or other flash memory. Additionally, or in the alternative, the medium may be a system or cache memory, or buffer space.

Design process 1610 may include hardware and software modules for processing a variety of input data structure types including netlist 1680. Such data structure types may reside, for example, within library elements 1630 and include a set of commonly used elements, circuits, and devices, including models, layouts, and symbolic representations, for a given manufacturing technology (e.g., different technology nodes, 32 nm, 45 nm, 90 nm, etc.). The data structure types may further include design specifications 1640, characterization data 1650, verification data 1660, design rules 1670, and test

data files **1685** which may include input test patterns, output test results, and other testing information. Design process **1610** may further include, for example, standard mechanical design processes such as stress analysis, thermal analysis, mechanical event simulation, process simulation for operations such as casting, molding, and die press forming, etc. One of ordinary skill in the art of mechanical design can appreciate the extent of possible mechanical design tools and applications used in design process **1610** without deviating from the scope and spirit of the invention. Design process **1610** may also include modules for performing standard circuit design processes such as timing analysis, verification, design rule checking, place and route operations, etc.

Design process **1610** employs and incorporates logic and physical design tools such as HDL compilers and simulation model build tools to process design structure **1620** together with some or all of the depicted supporting data structures along with any additional mechanical design or data (if applicable), to generate a second design structure **1690**. Design structure **1690** resides on a storage medium or programmable gate array in a data format used for the exchange of data of mechanical devices and structures (e.g., information stored in a IGES, DXF, Parasolid XT, JT, DRG, or any other suitable format for storing or rendering such mechanical design structures). Similar to design structure **1620**, design structure **1690** preferably comprises one or more files, data structures, or other computer-encoded data or instructions that reside on transmission or data storage media and that when processed by an ECAD system generate a logically or otherwise functionally equivalent form of one or more of the embodiments of the invention shown in FIGS. **1-3, 5** and **11**. In one embodiment, design structure **1690** may comprise a compiled, executable HDL simulation model that functionally simulates the devices shown in FIGS. **1-3, 5** and **11**.

Design structure **1690** may also employ a data format used for the exchange of layout data of integrated circuits and/or symbolic data format (e.g., information stored in a GDSII (GDS2), GL1, OASIS, map files, or any other suitable format for storing such design data structures). Design structure **1690** may comprise information such as, for example, symbolic data, map files, test data files, design content files, manufacturing data, layout parameters, wires, levels of metal, vias, shapes, data for routing through the manufacturing line, and any other data required by a manufacturer or other designer/developer to produce a device or structure as described above and shown in FIGS. **1-3, 5** and **11**. Design structure **1690** may then proceed to a stage **1695** where, for example, design structure **1690**: proceeds to tape-out, is released to manufacturing, is released to a mask house, is sent to another design house, is sent back to the customer, etc.

As has been described, a coherent attached processor proxy (CAPP) includes transport logic having a first interface configured to support communication with a system fabric of a primary coherent system and a second interface configured to support communication with an attached processor (AP) that is external to the primary coherent system and that includes a cache memory that holds copies of memory blocks belonging to a coherent address space of the primary coherent system. The CAPP further includes one or more master machines that initiate memory access requests on the system fabric of the primary coherent system on behalf of the AP, one or more snoop machines that service requests snooped on the system fabric, and a CAPP directory having a precise directory having a plurality of entries each associated with a smaller data granule and a coarse directory having a plurality of entries each associated with a larger data granule.

While various embodiments have been particularly shown as described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes

in form and detail may be made therein without departing from the spirit and scope of the claims. For example, although aspects have been described with respect to a computer system executing program code that directs the functions of the present invention, it should be understood that present invention may alternatively be implemented as a program product including a computer-readable storage device (e.g., volatile or non-volatile memory, optical or magnetic disk or other statutory manufacture) that stores program code that can be processed by a data processing system. Further, the term “coupled” as used herein is defined to encompass embodiments employing a direct electrical connection between coupled elements or blocks, as well as embodiments employing an indirect electrical connection between coupled elements or blocks achieved using one or more intervening elements or blocks. In addition, the term “exemplary” is defined herein as meaning one example of a feature, not necessarily the best or preferred example.

What is claimed is:

1. A method of data processing, comprising:

at a coherent attached processor proxy (CAPP) of a primary coherent system that serves as a participant in coherent communication within the primary coherent system on behalf of an attached processor (AP) external to the primary coherent system, maintaining, in a CAPP directory, entries corresponding to cachelines of interest to the AP that belong to a coherent address space of the primary coherent system, wherein the maintaining includes maintaining in a precise directory a plurality of entries each associated with a respective one of a plurality of smaller data granules and maintaining in a coarse directory a plurality of entries each associated with a respective one of a plurality of larger data granules;

the CAPP participating in coherent communication on a system fabric of the primary coherent system on behalf of the AP by reference to the coarse directory and the precise directory; and

the CAPP installing an entry in the coarse directory in a data-valid coherence state in advance of receipt of all of the larger data granule by the AP.

2. The method of claim 1, wherein:

the smaller data granule comprises a single cacheline of data; and

the larger data granule comprises a plurality of contiguous cachelines.

3. The method of claim 1, and further comprising:

in response to at least one AP command, the CAPP decomposes an entry in the coarse directory into a plurality of entries in the precise directory.

4. A method of data processing, comprising:

at a coherent attached processor proxy (CAPP) of a primary coherent system that serves as a participant in coherent communication within the primary coherent system on behalf of an attached processor (AP) external to the primary coherent system, maintaining, in a CAPP directory, entries corresponding to cachelines of interest to the AP that belong to a coherent address space of the primary coherent system, wherein the maintaining includes maintaining in a precise directory a plurality of entries each associated with a respective one of a plurality of smaller data granules and maintaining in a coarse directory a plurality of entries each associated with a respective one of a plurality of larger data granules;

the CAPP participating in coherent communication on a system fabric of the primary coherent system on behalf of the AP by reference to the coarse directory and the precise directory; and

the CAPP determining a composite coherence state for a target address by prioritizing a coherence state indicated by the precise directory over a coherence state indicated by the coarse directory.

5. A method of data processing, comprising: 5
at a coherent attached processor proxy (CAPP) of a primary coherent system that serves as a participant in coherent communication within the primary coherent system on behalf of an attached processor (AP) external to the primary coherent system, maintaining, in a CAPP 10 directory, entries corresponding to cachelines of interest to the AP that belong to a coherent address space of the primary coherent system, wherein the maintaining includes maintaining in a precise directory a plurality of entries each associated with a respective one of a plurality of smaller data granules and maintaining in a coarse 15 directory a plurality of entries each associated with a respective one of a plurality of larger data granules;
the CAPP participating in coherent communication on a system fabric of the primary coherent system on behalf 20 of the AP by reference to the coarse directory and the precise directory; and
in response to a target address of a snooped memory access request missing in the precise directory and hitting in the coarse directory, the CAPP providing a Retry response 25 and installing an entry corresponding to the target address in the precise directory.

* * * * *